

Dell PowerEdge RAID Controller (PERC) H310, H710, H710P, and H810 User's Guide

Regulatory Model: UCSA-801, UCSB-800, UCSM-800, UCPA-801,
UCPB-800, UCPM-800, and UCPE-800



Notes, Cautions, and Warnings



NOTE: A NOTE indicates important information that helps you make better use of your computer.



CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.



WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2013 Dell Inc. All Rights Reserved.

Trademarks used in this text: Dell™, the Dell logo, Dell Boomi™, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™, Venue™ and Vostro™ are trademarks of Dell Inc. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, vMotion®, vCenter®, vCenter SRM™ and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation.

2013 - 03

Rev. A02

Contents

1 Overview.....	9
Supported Operating Systems.....	10
Getting Help.....	10
Contacting Dell.....	11
Related Documentation.....	11
Documentation Feedback.....	11
2 Features.....	13
Physical Disk Power Management.....	13
Configured Spin Down Delay.....	13
Types Of Virtual Disk Initialization.....	14
Background Initialization Of Virtual Disks.....	14
Full Initialization Of Virtual Disks.....	14
Fast Initialization Of Virtual Disks.....	14
Consistency Checks.....	14
Disk Roaming.....	15
Using Disk Roaming.....	15
FastPath.....	15
Configuring FastPath-Capable Virtual Disks.....	15
Virtual Disk Migration.....	16
Migrating Virtual Disks.....	16
Virtual Disk Write Cache Policies.....	17
Write-Back And Write-Through.....	17
Conditions Under Which Write-Back Is Employed.....	17
Conditions Under Which Forced Write-Back With No Battery Is Employed.....	17
Virtual Disk Read Cache Policies.....	18
Reconfiguration Of Virtual Disks.....	18
Fault Tolerance.....	20
The SMART Feature.....	20
Automatic Replace Member With Predicted Failure.....	20
Patrol Read.....	21
Redundant Path Support (For PERC H810 Only).....	21
Physical Disk Failure Detection.....	22
Using Persistent Hot Spare Slots.....	22
Physical Disk Hot Swapping.....	22
Using Replace Member And Revertible Hot Spares.....	22
Controller Cache Preservation.....	23
Cache Preservation With Non-Volatile Cache (NVC).....	23

Recovering Cache Data.....	23
Battery Transparent Learn Cycle.....	23
TLC Time Frame.....	23
Conditions For Replacing The Battery.....	23
3 Deploying The PERC Card.....	25
Removing The PERC Controller.....	25
Installing The PERC Controller.....	27
Support For Internal Multiple Controllers.....	28
Setting Up Redundant Path Support On The PERC H810 Adapter.....	28
Reverting To Single Path Support From Redundant Path Support For PERC H810.....	28
4 Driver Installation.....	31
Pre-Installation Requirements For Windows Driver Installation.....	31
Creating The Device Driver Media For Windows Driver Installation.....	31
Downloading Drivers From The Dell Systems Service And Diagnostic Tools Media For Windows.....	31
Downloading Drivers From The Dell Support Website For Windows.....	32
Installing Driver During a Windows Server 2008/2008 R2 Installation.....	32
Installing Windows Server 2008/2008 R2 For A New RAID Controller.....	32
Updating Existing Windows Server 2008 Or Windows Server 2008 R2.....	33
Updating The Linux Driver.....	33
Installing Or Updating The RPM Driver Package With DKMS Support.....	34
Installing Or Updating The RPM Driver Package With KMOD Support.....	34
Installing Or Updating The RPM Driver Package With KMP Support.....	34
5 Management Applications For PERC Cards.....	37
Dell OpenManage Storage Management.....	37
BIOS Configuration Utility.....	37
Entering The BIOS Configuration Utility.....	37
Exiting The Configuration Utility.....	38
Menu Navigation Controls.....	38
Setting Up Virtual Disks.....	39
Virtual Disk Management.....	41
Creating Virtual Disks.....	41
Selecting Virtual Disk Parameters.....	42
Converting Physical Disk To RAID Capable For PERC H310.....	42
Converting Physical Disk To Non-RAID For PERC H310.....	43
Initializing Virtual Disks.....	43
Checking Data Consistency.....	43
Running A Data Consistency Check.....	43
Importing Or Clearing Foreign Configurations Using The VD Mgmt Menu.....	44
Importing Or Clearing Foreign Configurations Using The Foreign Configuration View Screen.....	44

Break Mirror.....	46
Managing Preserved Cache.....	46
Managing Dedicated Hot Spares.....	47
Deleting Virtual Disks.....	48
Deleting Disk Groups.....	48
Clearing The Configuration.....	48
BIOS Configuration Utility Menu Options.....	49
Virtual Disk Management (VD Mgmt).....	49
Virtual Disk Actions.....	51
Physical Disk Management (PD Mgmt).....	52
Physical Disk Actions.....	52
Rebuild.....	53
Controller Management (Ctrl Mgmt).....	53
Controller Management Actions.....	53
Foreign Configuration View.....	54
Physical Disk Management.....	54
Setting LED Blinking.....	54
Creating Global Hot Spares.....	55
Removing Global Or Dedicated Hot Spares.....	55
Replacing An Online Physical Disk.....	55
Restrictions and Limitations.....	56
Stopping Background Initialization.....	56
Performing A Manual Rebuild Of An Individual Physical Disk.....	56
Controller Management.....	57
Enabling Boot Support.....	57
Enabling Boot Support For A BIOS-Enabled Controller.....	57
Enabling BIOS Stop On Error.....	57
Enabling Auto Import.....	57
Disabling Auto Import.....	58
Restoring Factory Default Settings.....	58
UEFI RAID Configuration Utility.....	58
Entering The UEFI RAID Configuration Utility.....	58
Exiting The UEFI RAID Configuration Utility.....	59
Configuration Options.....	59
Controller Management Menu.....	59
Virtual Disk Management.....	60
Physical Disk Management Menu.....	60
Enclosure Management.....	60
6 CacheCade.....	61
CacheCade Virtual Disk Characteristics.....	61
Configuring And Managing CacheCade Virtual Disks.....	61

CacheCade Virtual Disk Management.....	62
Creating CacheCade Virtual Disks.....	62
Deleting CacheCade Virtual Disks.....	63
Reconfiguring CacheCade Virtual Disks.....	63
Automatic Reconfiguration Of CacheCade Virtual Disks.....	63
Manual Resizing Of CacheCade Virtual Disks.....	63
7 Security Key And RAID Management.....	65
Security Key Implementation.....	65
Security Key Management In The BIOS Configuration Utility.....	65
Local Key Management (LKM).....	65
Creating A Security Key.....	66
Changing The Security Key.....	66
Deleting A Security Key.....	67
Creating Secured Virtual Disks.....	67
Securing Pre-Existing Virtual Disks.....	67
Securing Pre-Existing Virtual Disks.....	68
Importing Or Clearing Secured Foreign Configurations And Secure Disk Migration.....	68
Instant Secure Erase.....	69
Troubleshooting Security Key Errors.....	69
Secured Foreign Import Errors.....	69
Failure to Select Or Configure Non Self-Encrypting Disks (Non-SED).....	69
Failure To Delete Security Key.....	70
Failure To Instant Secure Erase Task On Physical Disks.....	70
8 Troubleshooting.....	71
BIOS Error Messages.....	71
Discovery Error Message.....	71
Extra Enclosure Error Message.....	71
Cache Data Lost Error Message.....	71
Missing Disks In Virtual Disk Error Message.....	72
Previous Configuration Of Disks Removed Error Message.....	72
Missing Virtual Disks Error Message.....	72
Dirty Cache Data Error Message.....	73
BIOS Disabled Error Message.....	73
Drive Configuration Changes Error Message.....	73
Adapter At Baseport Not Responding Error Message.....	74
Offline Or Missing Virtual Drives With Preserved Cache Error Message.....	74
Virtual Disks Offline Error Message.....	74
Virtual Disks Degraded Error Message.....	74
Virtual Disks Partially Degraded Error Message.....	75
Memory Or Battery Problem Error Message.....	75





Firmware Fault State Error Message.....	75
Foreign Configuration Found Error Message.....	75
Foreign Configuration Not Found In <Ctrl> <R> Error Message.....	76
Previous Configuration Cleared Or Missing Error Message.....	76
Invalid SAS Topology Detected Error Message.....	76
Multibit ECC Errors Detected Error Messages.....	76
Configured Disks Removed Or Not Accessible Error Message.....	77
Battery Discharged Or Disconnected Error Message.....	77
Degraded State Of Virtual Disks.....	77
Memory Errors.....	78
Preserved Cache State.....	78
General Issues.....	78
PERC Card Has Yellow Bang In Device Manager.....	78
PERC Card Not Seen In Device Manager.....	78
No Hard Drives Found Error Message During Microsoft Windows Server 2003 Installation.....	78
Physical Disk Issues.....	79
Physical Disk In Failed State.....	79
Unable to Rebuild A Fault Tolerant Virtual Disk.....	79
Fatal Error Or Data Corruption Reported.....	79
Physical Disk Displayed As Blocked.....	79
Multiple Disks Become Inaccessible.....	80
Rebuilding A Failed Physical Disk.....	80
Virtual Disk Fails During Rebuild Using A Global Hot Spare.....	80
Virtual Disk Fails During Rebuild Using A Dedicated Hot Spare.....	80
Physical Disk Fails During Reconstruction On Redundant Virtual Disk.....	81
Virtual Disk Fails Rebuild Using A Dedicated Hot Spare.....	81
Physical Disk Takes A Long Time To Rebuild.....	81
SMART Errors.....	81
Smart Error Detected On A Physical Disk In A Redundant Virtual Disk.....	81
Smart Error Detected On A Physical Disk In A Non-Redundant Virtual Disk.....	82
Replace Member Errors.....	82
Source Disk Fails During Replace Member Operation.....	82
Target Disk Fails.....	82
General Disk Fails.....	82
Linux Operating System Errors.....	82
Virtual Disk Policy Is Assumed As Write-Through Error Message.....	82
Driver Does Not Auto-Build Into New Kernel.....	83
Unable To Register SCSI Device Error Message.....	83
Disk Carrier LED Indicators.....	84

9 Appendix: RAID Description.....	85
Summary Of RAID Levels.....	85

RAID Terminology.....	85
Disk Striping.....	85
Disk Mirroring.....	86
Spanned RAID Levels.....	86
Parity Data.....	86

Overview

The Dell PowerEdge RAID Controller (PERC) H310, H710, H710P, and H810 family of storage controller cards has the following characteristics:

- Complies with serial-attached SCSI (SAS) 2.0 providing up to 6 Gb/sec throughput.
- Supports Dell-qualified serial-attached SCSI (SAS) hard drives, SATA hard drives, and solid-state drives (SSDs).
 -  **NOTE:** Mixing SAS and SATA drives within a virtual disk is not supported. Also, mixing hard drives and SSDs within a virtual disk is not supported.
 -  **NOTE:** Mixing disks of different speeds (7,200 rpm, 10,000 rpm, or 15,000 rpm) and bandwidth (3 Gbps or 6 Gbps) PCIe while maintaining the same drive type (SAS or SATA) and technology (HDD or SSD) is supported.
- Offers RAID control capabilities which include support for RAID levels 0, 1, 5, 6, 10, 50, and 60.
 -  **NOTE:** PERC H310 supports RAID 5 with limited performance and does not support RAID 6 and RAID 60.
- Provides reliability, high performance, and fault-tolerant disk subsystem management.
- Offers Non-RAID support for direct access to disk drives (PERC H310 only).
 -  **NOTE:** Operating systems can directly access Non-RAID hard drives. A Non-RAID hard drive is not fault-tolerant and cannot be recovered if it fails. Only the PERC H310 controller allows configuration of disk drives as Non-RAID.

The following table compares the hardware configurations for the PERC H310, H710, H710P, and H810 cards.

Table 1. PERC H310, H710, H710P, and H810 Cards Hardware Configurations

Specification	H310	H710	H710P	H810
RAID Levels	0, 1, 5, 10, 50	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60	0, 1, 5, 6, 10, 50, 60
Enclosures Per Port	Not Applicable	Not Applicable	Not Applicable	Up to four enclosures
Processor	Dell Adapter SAS RAID-on-Chip, 8-port with LSI 2008 chipset	Dell Adapter SAS RAID-on-Chip, 8-port with LSI 2008 chipset	Dell Adapter SAS RAID-on-Chip, 8-port with LSI 2008 chipset	Dell Adapter SAS RAID-on-Chip, 8-port with LSI 2008 chipset
Backup Battery Unit (BBU)	No	Yes	Yes	Yes
Non-Volatile Cache	Not Applicable	Yes	Yes	Yes
Cache Memory	Not Applicable	512 MB DDR3 800 Mhz	1 GB DDR3 1333 Mhz	1 GB DDR3 1333 Mhz
Cache Function	Not Applicable	Write Back, Write Through, Adaptive Read Ahead, No Read Ahead, Read Ahead	Write Back, Write Through, Adaptive Read Ahead, No Read Ahead, Read Ahead	Write Back, Write Through, Adaptive Read Ahead, No Read Ahead, Read Ahead

Specification	H310	H710	H710P	H810
Maximum number of spans per Disk Group	8	8	8	8
Maximum number of Virtual Disks per Disk Group	16	16	16	16
Online Capacity Expansion	Yes	Yes	Yes	Yes
Dedicated and Global Hot Spares	Yes	Yes	Yes	Yes
Hot Swap Devices Supported	Yes	Yes	Yes	Yes
Hardware XOR Engine	Yes	Yes	Yes	Yes
Redundant Path Support	Yes	Yes	Yes	Yes

Supported Operating Systems

The PERC H310, H710, and H710P cards support the following operating systems:

- Microsoft Windows Server 2012
- Microsoft Windows Server 2008 including Hyper-V virtualization
- Microsoft Windows Server 2008 R2 and later
- Red Hat Enterprise Linux version 5.8 and later (32-bit and 64-bit)
- Red Hat Enterprise Linux version 6.2 and later (64-bit)
- SUSE Linux Enterprise Server version 10 SP4 (64-bit)
- SUSE Linux Enterprise Server version 11 SP2 (64-bit)
- VMware ESX 4.1 and ESXi 4.1 Update 2 and later
- VMware ESXi 5.0 and later



NOTE: For the latest list of supported operating systems and driver installation instructions, see the system documentation at dell.com/support/manuals. For specific operating system service pack requirements, see the Drivers and Downloads section at dell.com/support/manuals.

Getting Help

Contacting Dell



NOTE: Dell provides several online and telephone-based support and service options. If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog. Availability varies by country and product, and some services may not be available in your area.

To contact Dell for sales, technical support, or customer-service issues:

1. Go to **dell.com/contactdell**.
2. Select your country or region from the interactive world map.
When you select a region, the countries for the selected regions are displayed.
3. Select the appropriate language under the country of your choice.
4. Select your business segment.
The main support page for the selected business segment is displayed.
5. Select the appropriate option depending on your requirement.

Related Documentation



NOTE: For all storage controllers and PCIe SSD documents, go to **dell.com/support/manuals**, then, **Choose from a list of all Dell products → Servers, Storage & Networking → Dell Adapters** .



NOTE: For all Dell OpenManage documents, see **dell.com/support/manuals**, then, **Choose from a list of all Dell products → Software, Electronics & Peripherals → Software → Enterprise System Management** .



NOTE: For all operating system documents, go to **dell.com/support/manuals**, then, **Choose from a list of all Dell products → Software, Electronics & Peripherals → Software → Operating System** .



NOTE: For all PowerEdge and PowerVault documentation, go to **dell.com/support/manuals** and enter the system service tag to get your system documentation.

- *Dell PowerEdge RAID Controller (PERC) H310, H710, H710P, and H810 User's Guide*—Discusses features, installation, management and troubleshooting of PERC cards.
- Product-specific Dell PowerVault Owner's Manual—Discusses features, repair, management of PowerVault system.
- Product-specific Dell PowerEdge Owner's Manual—Discusses features, repair, management of PowerEdge system.

Documentation Feedback

If you have feedback for this document, write to **documentation_feedback@dell.com**. Alternatively, you can click on the **Feedback** link in any of the Dell documentation pages, fill up the form, and click **Submit** to send your feedback.

Features

Some of the features discussed for PowerEdge RAID Controller (PERC) H310, H710, H710P, and H810 cards are:

- Physical Disk Power Management
- Types of Virtual Disk Initialization
- Consistency Checks
- Disk Roaming
- Fast Path
- Virtual Disk Migration
- Reconfiguration Of Virtual Disks
- Fault Tolerance
- Patrol Read

Physical Disk Power Management

Physical disk power management is a power saving feature of the PERC H310, H710, H710P, and H810 cards. The feature allows disks to be spun down based on disk configuration and I/O activity. The feature is supported on all rotating SAS and SATA disks and includes unconfigured, configured and ☐ hot-spare disks. The physical disk power management feature is disabled by default. The feature can be enabled in the Dell Open Manage Storage Management application using the **Unified Extensible Firmware Interface (UEFI) RAID Configuration** utility. For more information, see the Dell OpenManage documentation at dell.com/support/manuals.

There are four power saving modes available:


No Power Savings (<i>default mode</i>)	All power savings features are disabled.
Balanced Power Savings	Spin down is enabled only for unconfigured and hot spare disks.
Maximum Power Savings	Spin down is enabled for configured, unconfigured and hot spare disks.
Customized Power Savings	All power savings features are customizable. You can specify a Quality of Service window during which the configured disks are excluded from spin-down.

Configured Spin Down Delay



NOTE: The **Configured Spin Down Delay** option is not applicable for the **No Power Savings** mode.


The amount of time to wait before spinning down disks can be set using **Configured Spin Down Delay**. The minimum value of the timer is 30 minutes (default) and the maximum is one day. Disks are spun down automatically and spun up when accessed. All disks are spun up on reboot.

 **NOTE:** There is a delay to I/O operations when a configured disk is being spun up.

Types Of Virtual Disk Initialization


You can initialize the virtual disks as described in the following sections:


 **CAUTION:** The initializing virtual disks task erases the files and file systems while keeping the virtual disk configuration intact. Initializing a virtual disk destroys all data on the virtual disk.


 **NOTE:** The initialization operations mentioned here are not applicable for ☐ Non-RAID disks.

Background Initialization Of Virtual Disks

Background Initialization (BGI) is an automated process that writes the parity or mirror data on newly created virtual disks. BGI does not run on RAID 0 virtual disks. You can control the BGI rate in the Dell OpenManage storage management application. Any change in the BGI rate does not take effect until the next BGI run.

 **NOTE:** You cannot disable BGI permanently. If you cancel BGI, it automatically restarts within five minutes. For information on stopping BGI, see the topic [Stopping Background Initialization](#).

 **NOTE:** Unlike full or fast initialization of virtual disks, background initialization does not clear data from the physical disks.


 **NOTE:** CC/BGI typically causes some loss in performance until the operation completes.

Consistency Check (CC) and BGI perform similar functions in that they both correct parity errors. However, CC reports data inconsistencies through an event notification, but BGI does not. You can start CC manually, but not BGI.

Full Initialization Of Virtual Disks

Performing a full initialization on a virtual disk overwrites all blocks and destroys any data that previously existed on the virtual disk. Full initialization of a virtual disk eliminates the need for the virtual disk to undergo a BGI. Full initialization can be performed after the creation of a virtual disk.

During full initialization, the host is not able to access the virtual disk. You can start a full initialization on a virtual disk by using the Slow Initialize option in the Dell OpenManage storage management application. For more information on using the **BIOS Configuration Utility** to perform a full initialization, see the topic [Initializing Virtual Disks](#).

 **NOTE:** If the system reboots during a full initialization, the operation aborts and a BGI begins on the virtual disk.

Fast Initialization Of Virtual Disks

A fast initialization on a virtual disk overwrites the first and last 8 MB of the virtual disk, clearing any boot records or partition information. The operation takes only 2–3 seconds to complete and is recommended when you are recreating virtual disks. To perform a fast initialization using the **BIOS Configuration Utility**, see the topic [Initializing Virtual Disks](#)

Consistency Checks

Consistency Check (CC) is a background operation that verifies and corrects the mirror or parity data for fault tolerant virtual disks. It is recommended that you periodically run a consistency check on virtual disks.

You can manually start a CC using the **BIOS Configuration Utility** or the Dell OpenManage storage management application. You can schedule CC to run on virtual disks using a Dell OpenManage storage management application. To start a CC using the **BIOS Configuration Utility**, see the topic [Checking Data Consistency](#).

Disk Roaming

Disk roaming is moving the physical disks from one cable connection or backplane slot to another on the same controller. The controller automatically recognizes the relocated physical disks and logically places them in the virtual disks that are part of the disk group. You can perform disk roaming only when the system is turned off.



CAUTION: Do not attempt disk roaming during RAID level migration (RLM) or online capacity expansion (OCE). This causes loss of the virtual disk.

Using Disk Roaming

Perform the following steps to use disk roaming:

1. Turn off the power to the system, physical disks, enclosures, and system components.
2. Disconnect power cables from the system.
3. Move the physical disks to desired positions on the backplane or the enclosure.
4. Perform a safety check. Make sure the physical disks are inserted properly.
5. Turn on the system.
The controller detects the RAID configuration from the configuration data on the physical disks.

FastPath

FastPath is a feature that improves application performance by delivering high I/O per second (IOPs). The Dell PowerEdge RAID Controller (PERC) H710P and H810 cards support **FastPath**.

FastPath is a further enhancement of the **Cut Through IO** (CTIO) feature, introduced in PERC H700 and PERC H800, to accelerate IO performance by reducing the IO processing overhead of the firmware. **CTIO** reduces the instruction count required to process a given IO. It also ensures that the optimal IO code path is placed close to the processor to allow faster access when processing the IO.

Under specific conditions with **FastPath**, the IO by-passes the controller cache and is committed directly to the physical disk from the host memory, through the second core of the dual-core RAID-on-Chip (ROC) on the controller. **FastPath** and **CTIO** are both ideal for random workloads with small blocks.



NOTE: The PERC H310 and PERC H710 do not support **FastPath**.

Both **CTIO** and **FastPath** provide enhanced performance benefits to SSD volumes, as they can fully capitalize on the lower access times and latencies of these volumes.

FastPath provides IO performance benefits to rotational HDD-based volumes configured with **Write Through** and **No Read Ahead** cache policies, specifically for read operations across all RAID levels and write operations for RAID 0.

Configuring FastPath-Capable Virtual Disks

All simple virtual disks configured with write cache policy **Write Through** and read cache policy **No Read Ahead** can utilize **FastPath**. Only IO block sizes smaller than virtual disk's stripe size are eligible for **FastPath**. In addition, there should be no background operations (rebuild, initialization) running on the virtual disks. **FastPath** will not be used if these operations are active.



NOTE: RAID 10, RAID 50, and RAID 60 virtual disks cannot use **FastPath**.

The following table summarizes the **FastPath**-eligibility of read and write IOs across the supported RAID levels.

Table 2. FastPath Eligibility Across Supported RAID Levels

	RAID 0	RAID 1	RAID 5	RAID 6
Read	Yes	Yes (Optimal and Degraded)	Yes (Optimal and Degraded)	Yes (Optimal and Degraded)
Write	Yes	No	No	No


In addition, FastPath-capable virtual disks are characterized by the following:

- The presence of CacheCade virtual disks on a controller disables the FastPath capability of all eligible HDD virtual disks. Eligible SSD virtual disks are still FastPath-capable.
- Secured virtual disks are not FastPath-capable.
- The **Physical Disk Power Management** feature is not applicable to FastPath-capable virtual disks.
- The addition or removal of a redundant path does not affect the FastPath-capability of virtual disks.


Virtual Disk Migration


The PERC H710, H710P, and H810 cards support migration of virtual disks from one controller to another without taking the target controller offline. The controller can import RAID virtual disks in optimal, degraded, or partially degraded states. You cannot import a virtual disk that is in an offline state. Disk migration pointers:

- Supports migration of VDs from PERC H700 and H800 to PERC H710P and H810
- Supports migration of volumes created within H710, H710P, or H810
- Supports migration of volumes created on H310 to H710, H710P, or H810
- Does not support migration from H700 or H800 to H310
- Does not support migration from H710, H710P, or H810 to H310
- Does not support backward migration from H310, H710, H710P, H810 to H800 and H700

 **NOTE:** The source controller must be offline prior to performing the disk migration.

 **NOTE:** Disks cannot be migrated to older revisions or generations of the PERC cards.

 **NOTE:** Non-RAID disks are supported only on the PERC H310 controller. Migration to any other PERC product is not supported.

 **NOTE:** Importing secured virtual disks is supported as long as the appropriate key (LKM) is supplied or configured.


When a controller detects a physical disk with an existing configuration, it flags the physical disk as foreign, and generates an alert indicating that a foreign disk was detected.

 **CAUTION:** Do not attempt disk roaming during RLM or OCE. This causes loss of the virtual disk.

Migrating Virtual Disks

To migrate virtual disks from PERC H700 or H800 to PERC H710, H710P, or H810:

1. Turn off the system.
2. Move the appropriate physical disks from the PERC H700 or H800 card to the PERC H710, H710P, or H810 card.

 **NOTE:** If you are replacing your PERC H700 or H800 with a PERC H710, H710P, or H810 card, see the *Owner's Manual* of your system at support.dell.com/manuals.

3. Boot the system and import the foreign configuration that is detected. You can do one of the following:

- Press <F> to automatically import the foreign configuration.
- Enter the **BIOS Configuration Utility** and navigate to the **Foreign Configuration View**.

 **NOTE:** For more information on accessing the **BIOS Configuration Utility**, see the topic [Entering The BIOS Configuration Utility](#).


 **NOTE:** For more information on **Foreign Configuration View**, see the topic [Foreign Configuration View](#).

4. Exit the **BIOS Configuration Utility** and reboot the system.

5. Ensure that all the latest drivers for the PERC H710, H710P, or H810 card (available at support.dell.com) are installed.


For more information, see the topic [Driver Installation](#).

Virtual Disk Write Cache Policies

 **NOTE:** PERC H310 only supports **Write-Through** mode.

The write cache policy of a virtual disk determines how the controller handles writes to the virtual disk. **Write-Back** and **Write-Through** are the two write cache policies that can be set on virtual disks individually.

All RAID volumes are presented as **Write-Through** to the operating system (Windows and Linux) independent of the actual write cache policy of the virtual disk. The PERC cards manage the data in cache independently of the operating system or any applications.


 **NOTE:** Use the Dell OpenManage storage management application or the **BIOS Configuration Utility** to view and manage virtual disk cache settings.


Write-Back And Write-Through

In **Write-Through** caching, the controller sends a data transfer completion signal to the host system when the disk subsystem has received all the data in a transaction.

In **Write-Back** caching, the controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction. The controller then writes the cached data to the storage device in the background.

The risk of using **Write-Back** cache is that the cached data can be lost if there is a power failure before it is written to the storage device. This risk is mitigated by using a Non-Volatile Cache on the PERC H710, H710P or H810 card.

 **NOTE:** The default cache setting for virtual disks is **Write-Back** caching.

 **NOTE:** Certain data patterns and configurations perform better with a **Write-Through** cache policy.

Conditions Under Which Write-Back Is Employed

Write-Back caching is used under all conditions in which the battery is present and in good condition.

Conditions Under Which Forced Write-Back With No Battery Is Employed


 **CAUTION:** It is recommended that you use a power backup system when forcing **Write-Back** to ensure there is no loss of data if the system suddenly loses power.

Write-Back mode is available when you select **Force WB with no battery**. When **Forced Write-Back** mode is selected, the virtual disk is in **Write-Back** mode even if the battery is not present.

Virtual Disk Read Cache Policies


The read policy of a virtual disk determines how the controller handles reads to that virtual disk. The read policies are:


- **Always Read Ahead**—Allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data is required soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
- **No Read Ahead**—Disables the **Read-Ahead** capability.
- **Adaptive Read Ahead**—Begins using **Read-Ahead** if the two most recent disk accesses occurred in sequential sectors. If the read requests are random, the controller reverts to **No Read Ahead** mode.

 **NOTE:** The default read cache setting for virtual disks is **Adaptive Read Ahead**.

Reconfiguration Of Virtual Disks

An online virtual disk can be reconfigured in ways that expands its capacity and/or change its RAID level.

 **NOTE:** Spanned virtual disks such as RAID 10, 50, and 60 cannot be reconfigured.


 **NOTE:** Reconfiguring Virtual Disks typically impacts disk performance until the reconfiguration operation is complete.

Online Capacity Expansion (OCE) can be done in two ways:

- If there is a single virtual disk in a disk group and free space is available, the virtual disk's capacity can be expanded within that free space.
- If a virtual disk is created and it does not use the maximum size of the disk group, free space is available.

Free space is also available when a disk group's physical disks are replaced by larger disks using the **Replace Member** feature. A virtual disk's capacity can also be expanded by performing an OCE operation to add more physical disks.

RAID Level Migration (RLM) refers to changing a virtual disk's RAID level. Both RLM and OCE can be done at the same time so that a virtual disk can simultaneously have its RAID level changed and its capacity increased. When a RLM/OCE operation is complete, a reboot is not required. See the following table for a list of RLM/OCE possibilities. The source RAID level column indicates the virtual disk RAID level before the RLM/OCE and the target RAID level column indicates the RAID level after the operation has completed.

 **NOTE:** If the controller already contains the maximum number of virtual disks, you cannot perform a RAID level migration or capacity expansion on any virtual disk.



 **NOTE:** The controller changes the write cache policy of all virtual disks undergoing a RLM/OCE to **Write-Through** until the RLM/OCE is complete.

Table 3. RAID Level Migration

Source RAID Level	Target RAID Level	Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansion Possible	Description
RAID 0	RAID 0	1	2 or more	Yes	Increases capacity by adding disks
RAID 0	RAID 1	1	2	No	Converts non-redundant virtual

Source RAID Level	Target RAID Level	Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansion Possible	Description
					disk into a mirrored virtual disk by adding one disk.
RAID 0	RAID 5	1 or more	3 or more	Yes	At least one disk needs to be added for distributed parity data.
RAID 0	RAID 6	1 or more	4 or more	Yes	At least two disks need to be added for dual distributed parity data.
RAID 1	RAID 0	2	2 or more	Yes	Removes redundancy while increasing capacity.
RAID 1	RAID 5	2	3 or more	Yes	Maintains redundancy while doubling capacity.
RAID 1	RAID 6	2	4 or more	Yes	Two disks required to be added for distributed parity data.
RAID 5	RAID 0	3 or more	3 or more	Yes	Converts to a non-redundant virtual disk and reclaims disk space used for distributed parity data.
RAID 5	RAID 5	3 or more	4 or more	Yes	Increases capacity by adding disks
RAID 5	RAID 6	3 or more	4 or more	Yes	At least one disk needs to be added for dual distributed parity data.
RAID 6	RAID 0	4 or more	4 or more	Yes	Converts to a non-redundant virtual disk and reclaims

Source RAID Level	Target RAID Level	Number of Physical Disks (Beginning)	Number of Physical Disks (End)	Capacity Expansion Possible	Description
					disk space used for distributed parity data.
RAID 6	RAID 5	4 or more	4 or more	Yes	Removes one set of parity data and reclaims disk space used for it.
RAID 6	RAID 6	4 or more	5 or more	Yes	Increases capacity by adding disks.

 **NOTE:** The total number of physical disks in a disk group cannot exceed 32. You cannot perform RAID level migration and expansion on RAID levels 10, 50, and 60.

Fault Tolerance

The list of features of the PERC cards that provide fault tolerance to prevent data loss is as follows:


- Support for Self Monitoring and Reporting Technology (SMART)
- Support for Patrol Read
- Redundant path support (for PERC H810 only)
- Physical disk failure detection
- Physical disk rebuild using hot spares
- Controller Cache Preservation
- Battery and Non-Volatile Cache backup of controller cache to protect data
- Detection of batteries with low charge after boot up

The next sections describe some methods to achieve fault tolerance.

The SMART Feature

The SMART feature monitors certain physical aspects of all motors, heads, and physical disk electronics to help detect predictable physical disk failures. SMART-compliant physical disks have attributes for which data can be monitored to identify changes in values and determine whether the values are within threshold limits. Many mechanical and electrical failures display some degradation in performance before failure.

A SMART failure is also referred to as a predicted failure. There are numerous factors that relate to predicted physical disk failures, such as a bearing failure, a broken read/write head, and changes in spin-up rate. In addition, there are factors related to read/write surface failure, such as seek error rate and excessive bad sectors.


 **NOTE:** For detailed information on SCSI interface specifications, see t10.org and for detailed information on SATA interface specifications, see t13.org.

Automatic Replace Member With Predicted Failure

A **Replace Member** operation can occur when there is a SMART predictive failure reporting on a physical disk in a virtual disk. The automatic **Replace Member** is initiated when the first SMART error occurs on a physical disk that is part

of a virtual disk. The target disk needs to be a hot spare that qualifies as a rebuild disk. The physical disk with the SMART error is marked as **failed** only after the successful completion of the **Replace Member**. This avoids putting the array in degraded status.

If an automatic **Replace Member** occurs using a source disk that was originally a hot spare (that was used in a rebuild), and a new disk added for the **Replace Member** operation as the target disk, the hot spare reverts to the hot spare state after a successful **Replace Member** operation.


 **NOTE:** To enable the automatic **Replace Member**, use the Dell OpenManage storage management application. For more information on automatic **Replace Member**, see the topic [Dell OpenManage Storage Management](#). For information on manual **Replace Member**, see the topic [Replacing An Online Physical Disk](#).

Patrol Read

The **Patrol Read** feature is designed as a preventative measure to ensure physical disk health and data integrity. **Patrol Read** scans for and resolves potential problems on configured physical disks. The Dell OpenManage storage management application can be used to start **Patrol Read** and change its behavior.

The following is an overview of **Patrol Read** behavior:

- **Patrol Read** runs on all disks on the controller that are configured as part of a virtual disk, including hot spares.
- **Patrol Read** does not run on physical disks that are not part of a virtual disk or are in **Ready** state.
- **Patrol Read** adjusts the amount of controller resources dedicated to **Patrol Read** operations based on outstanding disk I/O. For example, if the system is busy processing I/O operation, then **Patrol Read** uses fewer resources to allow the I/O to take a higher priority.
- **Patrol Read** does not run on any disks involved in any of the following operations:
 - Rebuild
 - Replace Member
 - Full or Background Initialization
 - CC
 - RLM or OCE

 **NOTE:** By default, **Patrol Read** automatically runs every seven days on configured SAS and SATA hard drives. **Patrol Read** is not necessary on SSD and is disabled by default.


For more information on **Patrol Read**, see the Dell OpenManage documentation at dell.com/support/manuals.

Redundant Path Support (For PERC H810 Only)

The PERC H810 adapter can detect and use redundant paths to disks contained in enclosures. This provides the ability to connect two SAS cables between a controller and an enclosure for path redundancy. The controller is able to tolerate the failure of a cable or Enclosure Management Module (EMM) by utilizing the remaining path.

When redundant paths exist, the controller automatically balances I/O load through both paths to each disk. Load balancing increases throughput to virtual disks in storage enclosures and is automatically turned on when redundant paths are detected. The ability to load balance I/O can be disabled using the Dell OpenManage storage management application.


To set up your hardware to support redundant paths, see the topic [Setting Up Redundant Path Support On The PERC H810 Adapter](#).

 **NOTE:** This support for redundant paths refers to path-redundancy only and not to controller-redundancy.

Physical Disk Failure Detection

Failed physical disks are detected and rebuilds automatically start to new disks that are inserted into the same slot. Automatic rebuilds can also happen transparently with hot spares. If you have configured hot spares, the controllers automatically try to use them to rebuild failed physical disks.

Using Persistent Hot Spare Slots


 **NOTE:** The persistent hot spare slot feature is disabled by default.

The PERC H310, H710, H710P, and H810 cards can be configured so that the system backplane or storage enclosure disk slots are dedicated as hot spare slots. This feature can be enabled using the Dell OpenManage storage management application.

Once enabled, any slots with hot spares configured automatically become persistent hot spare slots. If a hot spare disk fails or is removed, a replacement disk that is inserted into the same slot automatically becomes a hot spare with the same properties as the one it is replacing. If the replacement disk does not match the disk protocol and technology, it does not become a hot spare.

For more information on persistent hot spares, see the Dell OpenManage documentation at dell.com/support/manuals.

Physical Disk Hot Swapping


 **NOTE:** To check if the backplane supports hot swapping, see the *Owner's Manual* of your system.


Hot swapping is the manual replacement of a disk while the PERC H310, H710, H710P, or H810 cards are online and performing their normal functions. The following requirements must be met before hot swapping a physical disk:

- The system backplane or enclosure must support hot swapping for the PERC H310, H710, H710P or H810 cards to support hot swapping.
- The replacement disk must be of the same protocol and disk technology. For example, only a SAS hard drive can replace a SAS hard drive; only a SATA SSD can replace a SATA SSD.
- The replacement disk must be of equal or greater capacity than the one it is replacing.

Using Replace Member And Revertible Hot Spares

The **Replace Member** functionality allows a previously commissioned hot spare to be reverted to a usable hot spare. When a disk failure occurs within a virtual disk, an assigned hot spare (dedicated or global) is commissioned and begins rebuilding until the virtual disk is optimal. After the failed disk is replaced (in the same slot) and the rebuild to the hot spare is complete, the controller automatically starts to copy data from the commissioned hot spare to the newly-inserted disk. After the data is copied, the new disk is a part of the virtual disk and the hot spare is reverted to being a ready hot spare. This allows hot spares to remain in specific enclosure slots. While the controller is reverting the hot spare, the virtual disk remains optimal.

 **NOTE:** The controller automatically reverts a hot spare only if the failed disk is replaced with a new disk in the same slot. If the new disk is not placed in the same slot, a manual **Replace Member** operation can be used to revert a previously commissioned hot spare.

 **NOTE:** A **Replace Member** operation typically causes a temporary impact to disk performance. Once the operation completes, performance returns to normal.

Controller Cache Preservation

The controller is capable of preserving its cache in the event of a system power outage or improper system shutdown. The PERC H710, H710P, and H810 controllers are attached to a Battery Backup Unit (BBU) that provides backup power during system power loss to preserve the controller's cache data.

Cache Preservation With Non-Volatile Cache (NVC)

In essence, the NVC module allows controller cache data to be stored indefinitely. If the controller has data in the cache memory during a power outage or improper system shutdown, a small amount of power from the battery is used to transfer cache data to a non-volatile flash storage where it remains until power is restored and the system is booted.


Recovering Cache Data

The dirty cache LED that is located on the H710 and H810 cards can be used to determine if cache data is being preserved.

If a system power loss or improper system shutdown has occurred:


1. Restore the system power.
2. Boot the system.
3. To enter the **BIOS Configuration Utility**, select **Managed Preserved Cache** in the controller menu.
If there are no virtual disks listed, all preserved cache data has been written to disk successfully.

Battery Transparent Learn Cycle

 **NOTE:** Batteries are only supported on PERC H710, H710P, and H810 cards.

A transparent learn cycle is a periodic operation that calculates the charge that is remaining in the battery to ensure there is sufficient energy. The operation runs automatically, and causes no impact to the system or controller performance.

The controller automatically performs the Transparent Learn Cycle (TLC) on the battery to calibrate and gauge its charge capacity once every 90 days. The operation can be performed manually, if required.

 **NOTE:** Virtual disks stay in **Write Back** mode, if enabled, during transparent learn cycle. When the TLC completes, the controller sets the next TLC to +90 days.


TLC Time Frame


The time frame for completion of a learn cycle is a function of the battery charge capacity and the discharge and charge currents used. For PERC H710 or H810 cards, the expected time frame for completion of a learn cycle is approximately seven hours.


Conditions For Replacing The Battery

The PERC battery is marked **Failed** when the state or health of the battery is declared bad. If battery is declared failed then the firmware runs learn cycles in subsequent reboots until the battery is replaced. The virtual disk then transitions to **Write Back** mode.

Deploying The PERC Card

 **CAUTION:** Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the system.

 **NOTE:** For information on removing and reinstalling system parts, see the Owner's Manual of the system at dell.com/support/manuals.

 **NOTE:** For more information on your storage controller, see the relevant storage controller documentation at dell.com/support/manuals.

This document provides a set of high level installation and removal instructions for the following Dell PowerEdge RAID Controllers (PERC):

- PERC H310 Adapter
- PERC H310 Mini Monolithic
- PERC H310 Mini Blade
- PERC H710 Adapter
- PERC H710P Adapter
- PERC H710 Mini Monolithic
- PERC H710P Mini Monolithic
- PERC H710 Mini Blade
- PERC H710P Mini Blade
- PERC H810 Adapter

Removing The PERC Controller

To remove the PERC controller:

1. Perform a controlled shutdown of the system and attached peripherals.
2. Disconnect the system from the electrical outlet and remove the system cover.
3. Remove the storage controller from the PCIe slot of the system.

 **CAUTION:** When removing or replacing the PERC Mini Monolithic Controller, hold the card by its edges marked by the blue touch points. Do not handle the card while holding the battery or the heatsink.

For more information, see to the relevant figures below.




Figure 1. Removing and Installing the PERC Adapter

- 1. PCIe slot
- 2. SAS cable connectors (2)
- 3. PERC adapter
- 4. screw (1)




Figure 2. Removing and Installing the PERC Mini Blade Controller

1. release lever (2)
2. PERC stack-up connector
3. PERC mini blade controller
4. screws (2)




Figure 3. Removing and Installing the PERC Mini Monolithic Controller

- | | |
|-----------------------------------|--|
| 1. release lever (2) | 6. touch points (2) |
| 2. storage-controller card holder | 7. storage-controller card connector on the system board |
| 3. storage controller card | 8. guide pins (2) |
| 4. battery | |
| 5. heatsink | |

Installing The PERC Controller

To install the PERC controller:

1. Perform a controlled shutdown of the system and attached peripherals.
2. Disconnect the system from the electrical outlet and remove the system cover.

3. Install the storage controller in the appropriate controller slot and connect all the cables to the storage controller.



CAUTION: When removing or replacing the PERC Mini Monolithic Controller, hold the card by its edges marked by the blue touch points. Do not handle the card while holding the battery or the heatsink.

For more information, see the relevant figures in [Removing The PERC Controller](#).

4. Replace the system cover.
5. Reconnect the system to its electrical outlet and turn the system on, including any attached peripherals.

Support For Internal Multiple Controllers

On PowerEdge systems with dual bays (backplanes), the disk drives are managed by independent PERC H710P cards. The PERC H710P cards do not share the disks or the RAID volumes. A drive can be identified using the following as guideline: **Controller: Bay/Backplane: Slot**.

Setting Up Redundant Path Support On The PERC H810 Adapter

The PERC H810 card can detect and use redundant paths to disks contained in enclosures. With redundant paths to the same device, if one path fails, another path can be used to communicate between the controller and the device.

To set up a configuration with redundant paths, both ports on a controller must be cabled to the **In** ports of a single enclosure.

To add multiple enclosures, both **Out** ports of the first enclosure must be cabled to the **In** ports of the next enclosure.

If the connection between an **Out** port on the controller and an **In** port on an enclosure fails, an alternate path exists through the second **Out** port on the controller and the second **In** port on the enclosure. For more information, see the following figure.



NOTE: The PERC H810 card supports redundant paths when used with Dell PowerVault MD1200 and Dell PowerVault MD1220 disk storage enclosures.

Perform the following steps to configure the hardware to utilize redundant paths on the PERC H810 card:

1. Set up an enclosure on the PERC H810 card.
2. Connect two SAS cables from the **Out** ports on your PERC H810 card to the **In** ports of the external enclosure. For more information, see the following figure.



NOTE: For information on Unified Mode, see the enclosure documentation that was shipped with the enclosure.

3. To add multiple enclosures, cable both **Out** ports of the first enclosure to both **In** ports of the next enclosure. After you set up the hardware, the controller detects the redundant paths and automatically utilizes them to balance the I/O load.

Reverting To Single Path Support From Redundant Path Support For PERC H810

If you need to revert to single path support from redundant path support, shut down the system and remove the exact same cables that were added to support redundant path support, leaving only one connection between the controller and enclosures. After you remove the cable and power up the system, ensure that there are no warning messages during boot, and that all virtual disks are online and optimal. If you are using Dell OpenManage, see the Dell OpenManage documentation at dell.com/support/manuals for additional instructions.

CAUTION: If you remove any cables other than the ones added to enable redundant path support, the enclosure and disks can get disconnected, and virtual disk may fail.

Perform the following steps to configure the hardware to utilize redundant paths on the PERC H810 card:

1. Set up an enclosure on the PERC H810 card.
2. Connect two SAS cables from the **Out** ports on your PERC H810 card to the **In** ports of the external enclosure. For more information, see the following figure.

NOTE: For information on Unified Mode, see the enclosure documentation that was shipped with the enclosure.

3. To add multiple enclosures, cable both **Out** ports of the first enclosure to both **In** ports of the next enclosure.
4. After you set up the hardware, the controller detects the redundant paths and automatically utilizes them to balance the I/O load.




Figure 4. Redundant Path Support Configuration With Two Enclosures

1. server
2. storage

Driver Installation

The Dell PowerEdge RAID Controller (PERC) H310, H710, H710P, and H810 cards require software drivers to operate with the supported operating systems.

This chapter contains the procedures for installing the drivers for the PERC H310, H710, H710P, and H810 cards.



NOTE: For more information on VMware ESX drivers, see the VMware ESX documentation at dell.com/support/manuals.



NOTE: To check operating system compatibility, see dell.com/support/manuals.

The two methods for installing a driver discussed in this chapter are:

Installing a driver during operating system installation

Use this method if you are performing a new installation of the operating system and want to include the drivers.

Updating existing drivers

Use this method if the operating system and the PERC H310, H710, H710P, and H810 family of controllers are already installed and you want to update to the latest drivers.

Pre-Installation Requirements For Windows Driver Installation

Before you install the operating system:

- Read the Microsoft *Getting Started* document that shipped with your operating system.
- Ensure that your system has the latest BIOS, firmware, and driver updates. If required, download the latest BIOS, firmware, and driver updates from dell.com/support.
- Create a device driver media (diskette, USB drive, CD, or DVD).

Creating The Device Driver Media For Windows Driver Installation

Use one of the following two methods to create the device driver media:

- Downloading drivers from the Dell systems service and diagnostic tools media for Windows
- Downloading drivers from the Dell support website for Windows

Downloading Drivers From The Dell Systems Service And Diagnostic Tools Media For Windows

To download drivers from the Dell Systems Service and Diagnostic Tools media:

1. Insert the *Dell Systems Service and Diagnostics Tools* media in a system.
The **Welcome to Dell Service and Diagnostic Utilities** screen is displayed.
2. Select your system model and operating system (Microsoft Windows Server 2008).
3. Click **Continue**.

4. From the list of drivers displayed, select the driver you require. Select the self-extracting zip file and click **Run**. Copy the driver to a diskette drive, CD, DVD, or USB drive. Repeat this step for all the drivers you require.
5. During the operating system installation, use the media that you created with the **Load Driver** option to load mass storage drivers. For more information on reinstalling the operating system, see the relevant section for your operating system below.

Downloading Drivers From The Dell Support Website For Windows

To download drivers from the Dell Support website:

1. Go to **dell.com/support**.
2. Select your line of business.
3. In the **Popular support tools** section, click **Drivers and Downloads**.
4. Enter the service tag of your system in the **Choose by Service Tag to get started** field or select **Choose from a list of all Dell products**.
5. Select the **System Type, Operating System**, and **Category** from the drop-down list.
The drivers that are applicable to your selection are displayed.
6. Download the drivers that you require to a diskette drive, USB drive, CD, or DVD.
7. During the operating system installation, use the media that you created with the **Load Driver** option to load mass storage drivers. For more information on reinstalling the operating system, see the relevant section for your operating system below.

Installing Driver During a Windows Server 2008/2008 R2 Installation

To install the driver:

1. Boot the system using the Windows Server 2008, or Windows Server 2008 R2 media.
2. Follow the on-screen instructions until you reach **Where do you want to install Windows Server 2008** window and then select **Load driver**.
3. The system prompts you to insert the media. Insert the installation media and browse to the proper location.
4. Select the appropriate PERC H310, H710, H710P, or H810 card from the list, click **Next** and continue installation.


Installing Windows Server 2008/2008 R2 For A New RAID Controller

Perform the following steps to configure the driver for the RAID controller on a system that already has Windows installed:


1. Turn off the system.
2. Install the new RAID controller in the system.
For detailed instructions on installing and cabling the RAID controller in the system, see the topic [Deploying The PERC Card](#).
3. Turn on the system.
The **Found New Hardware Wizard** screen displays the detected hardware device.
4. Click **Next**.
5. On the **Locate device driver** screen, select **Search for a suitable driver for my device** and click **Next**.
6. Browse and select the drivers from the **Locate Driver Files** screen.
7. Click **Next**.
The wizard detects and installs the appropriate device drivers for the new RAID controller.
8. Click **Finish** to complete the installation.

9. Reboot the system when prompted.


Updating Existing Windows Server 2008 Or Windows Server 2008 R2

 **NOTE:** Close all applications on your system before you update the driver.


1. Insert the media (CD, DVD, or USB drive) containing the driver.
2. Select **Start → Settings → Control Panel → System**.
The **System Properties** screen is displayed.

 **NOTE:** The path to **System** might vary depending on the operating system family.


3. Click on the **Hardware** tab.
4. Click **Device Manager**.
The **Device Manager** screen is displayed.

 **NOTE:** The path to **Device Manager** might vary depending on the operating system family.


5. Expand **SCSI and RAID Controllers** by double-clicking the entry or by clicking on the plus symbol next to **SCSI and RAID Controller**.

 **NOTE:** In Windows 2008 and Windows 2008 R2, the PERC card is listed under **Storage Controllers**.


6. Double-click the RAID controller for which you want to update the driver.
7. Click the **Driver** tab and click **Update Driver**.
The screen to update the device driver wizard is displayed.
8. Select **Install from a list or specific location**.
9. Click **Next**.
10. Follow the steps in the wizard and browse to the location of the driver files.
11. Select the INF file from the driver media (CD, DVD, or other media).
12. Click **Next** and continue the installation steps in the wizard.
13. Click **Finish** to exit the wizard and reboot the system for the changes to take place.

 **NOTE:** Dell provides the Dell Update Package (DUP) to update drivers on systems running Windows Server 2008 or Windows Server 2008 R2 operating system. DUP is an executable application that updates drivers for specific devices. DUP supports command line interface and silent execution. For more information, see dell.com/support.

Updating The Linux Driver


 **NOTE:** PERC H310/H710/H710P/H810, PERC H700/H800 cards and both the PERC 5 and PERC 6 family of controllers use the same driver and do not require separate driver installations.

Use the procedures in this section to update the driver for Linux. To ensure that you have the current version of the driver, download the updated Linux driver from dell.com/support.

 **NOTE:** The driver update disk (DUD) images are created only for those operating system releases in which the native (in-box) driver is insufficient for installation. In the event that an operating system is being installed with a corresponding DUD image, follow the instructions below. If not, proceed with using the native device driver and then skip to the topic Installing or Updating the RPM Driver Package With DKMS Support.

Installing Or Updating The RPM Driver Package With DKMS Support

 **NOTE:** This procedure is applicable for Red Hat Enterprise Linux 5 SP7 and SUSE Enterprise Linux 10 SP4.

 **NOTE:** For SUSE Enterprise Linux 10 SP4, immediately following the operating system installation, download the latest driver from dell.com/support, and update the driver using the procedures detailed in this section.


Perform the following steps to install the RPM package with DKMS support:

1. Uncompress the gzipped tarball driver release package.
2. Install the DKMS package using the command: **rpm -ihv megaraid_sas- <version> .noarch.rpm.**
3. Install the driver package using the command: **rpm -ihv megaraid_sas- <version> .noarch.rpm.**

 **NOTE:** Use **rpm -Uvh <package name>** when updating an existing package.

4. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
5. Verify that the driver has been loaded with the following system commands: **modinfo megaraid_sas** and **dkms status**.

Installing Or Updating The RPM Driver Package With KMOD Support

 **NOTE:** This procedure is applicable for Red Hat Enterprise Linux 6 SP2.


Perform the following steps to install the RPM package with KMOD support:

1. Uncompress the gzipped tarball driver release package.
2. Install the driver package using the command: **rpm -ihv kmodmegaraid_sas-<version>.rpm.**

 **NOTE:** Use **rpm -Uvh <package name>** when upgrading an existing package.

3. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
4. Verify that the driver has been loaded with the following system commands: **modinfo megaraid_sas**.

Installing Or Updating The RPM Driver Package With KMP Support

 **NOTE:** This procedure is applicable for SUSE Enterprise Linux 11 SP2.

Perform the following steps to install the RPM package with KMP support:

1. Uncompress the gzipped tarball driver release package.
2. Install the driver package using the command: **rpm -ihv kmpmegaraid_sas- <version>.rpm.**

 **NOTE:** Use **rpm -Uvh <package name>** when updating an existing package.

3. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.
4. Verify that the driver has been loaded with the following system commands: **modinfo megaraid_sas**.

Upgrading The Kernel

When upgrading to a new kernel, you must reinstall the DKMS-enabled driver packages. Perform the following steps to update or install the driver for the new kernel:

1. In a terminal window, type the following: **dkms build -m <module_name> -v <module version> -k <kernel version>**
dkms install -m <module_name> -v <module version> -k <kernel version>.
2. To check whether the driver is successfully installed in the new kernel, type: **dkms status**.

You see a message similar to the following one: <driver name>, <driver version>, <new kernel version>: installed.

3. If the previous device driver is in use, you must reboot the system for the updated driver to take effect.

Management Applications For PERC Cards


Dell OpenManage Storage Management applications enable you to manage and configure the RAID system, create and manage multiple disk groups, control and monitor multiple RAID systems, and provide online maintenance. The management applications for PERC H310, H710, H710P, and H810 include:

- Dell OpenManage Storage Management
- BIOS Configuration Utility (<Ctrl> <R>)
- Unified Extensible Firmware Interface (UEFI) RAID Configuration Utility

 **NOTE:** Dell serial-attached SCSI (SAS) RAID Storage Manager is not supported for SED management.


Dell OpenManage Storage Management

The Dell OpenManage Storage Management is a storage management application for Dell systems that provides enhanced features for configuring a system's locally-attached RAID and Non-RAID disk storage. The Dell OpenManage storage management application enables you to perform controller and enclosure functions for all supported RAID controllers and enclosures from a single graphical or command-line interface without requiring the use of the controller BIOS utilities. The graphical user interface (GUI) is wizard-driven with features for novice and advanced users, and detailed online help. Using the Dell OpenManage storage management application, you can protect your data by configuring data-redundancy, assigning hot spares, or rebuilding failed physical disks. The command line interface available on selected operating systems to perform RAID management tasks is fully featured and scriptable.


 **NOTE:** For more information, see the *Dell OpenManage Storage Management User's Guide* at dell.com/support/manuals.

BIOS Configuration Utility

The **BIOS Configuration Utility**, also known as <Ctrl> <R>, is a storage management application embedded on the PERC H310, H710, H710P, and H810 cards that configures and maintains RAID disk groups and virtual disks. <Ctrl> <R> is independent of the operating system.

 **NOTE:** Use the **BIOS Configuration Utility (<Ctrl> <R>)** for initial setup and disaster recovery. You can use advanced features through Dell OpenManage storage management application and Dell SAS RAID storage manager.

The following sections provide information about using the **BIOS Configuration Utility (<Ctrl> <R>)**. For more information, see the online help option by pressing <F1> in the **BIOS Configuration Utility (<Ctrl> <R>)**.

 **NOTE:** The PERC PER H310, H710, H710P, and H810 card configuration utility refreshes the screen to show changes to information. The refresh occurs when you press <F5> or every 15 seconds.

Entering The BIOS Configuration Utility


Perform the following steps to enter the **BIOS Configuration Utility (<Ctrl> <R>)** when you boot the system:


1. Turn on the system.

A BIOS screen displays information about the controller and configuration.

2. During startup, press **<Ctrl> <R>** when prompted by the BIOS screen.
3. Use the arrow keys to select the RAID controller you want to configure, and press **<Enter>** to access the management menus for the controller.

If there is only one controller, the **Virtual Disk Management** screen for that controller is displayed. If there is more than one controller, the main menu screen is displayed. The screen lists the RAID controllers.

 **NOTE:** You can access multiple controllers through the **BIOS Configuration Utility** (**<Ctrl> <R>**) by pressing **<F12>**.

 **NOTE:** You can access PERC H700, H800, H310, H710, H710P, or H810 cards from the same BIOS if the PERC 6/H700/H800 firmware is 6.2.0-0013 or later.

Exiting The Configuration Utility

To exit the **BIOS Configuration Utility** (**<Ctrl> <R>**)

1. Press **<Esc>** on any menu screen.
If there is only one controller, then a dialog box is displayed to confirm your choice.
2. Select **OK** to exit and press **<Enter>**.
If multiple controllers are present, then the **<Esc>** key brings you to the **Controller Selection** screen.
3. Press **<Esc>** again to reach the exit screen.
A dialog box is displayed to confirm your choice.
4. Select **OK** to exit and press **<Enter>**.

Menu Navigation Controls

The following table displays the menu keys you can use to move between the different screens in the **BIOS Configuration Utility** (**<Ctrl> <R>**).

Table 4. Menu Navigation Keys

Notation	Meaning and Use	Example
right-arrow key	Use the right-arrow key to open a submenu, move from a menu heading to the first submenu, or move to the first item in that submenu. If you press the right-arrow key at a menu heading, the submenu expands. Press it again to go to the first item in the submenu. The right-arrow key is also used to close a menu list in a popup window. Word wrap is supported.	Start → Programs
left-arrow key	Use the left-arrow key to close a submenu, move from a menu item to the menu heading for that item, or move from a submenu to a higher level menu. If you press the left-arrow key at a menu heading, the submenu collapses. Press it again to go to the higher-level menu. Word wrap is supported.	Controller 0 ← Disk Group 1
up-arrow key	Use the up-arrow key to move to the upper menu items within a menu or to a higher level menu. You can also use the up-arrow key to close a menu list in a popup window, such as the stripe element size menu. Word wrap is supported.	Virtual Disk 1 ↑ Virtual Disk 4

Notation	Meaning and Use	Example
down-arrow key	Use the down-arrow key to move to the lower menu items within a menu or to a lower level menu. You can also use the down-arrow key to open a menu list in a popup window, such as the stripe element size menu, and select a setting. Word wrap is supported.	Virtual Disk 1 ↓ Virtual Disk 4
<Enter>	After you highlight a menu item, press <Enter> to select that item. An options menu for the menu item opens. It applies to only certain menu items, such as Virtual Disk # . In a list of options for that item, such as the write policy for a virtual disk, highlight a setting, such as Write-Through , and press <Enter> to select it.	Select Add New VD and press <Enter> to create a new virtual disk.
<Esc>	After you expand a pop-up window, press <Esc> to close the window. You can continue to press <Esc> to exit the BIOS Configuration Utility (<Ctrl> <R>).	Press <Esc> to return to the VD Mgmt screen.
<Tab>	Press <Tab> to move the cursor to the next control on a dialog box or page.	Press <Tab> to move the cursor to the next parameter you want to change.
<Shift> <Tab>	Press <Shift> <Tab> to move the cursor to the previous control on a dialog or page.	Press <Shift> <Tab> to move the cursor from Sort By to the previously selected PD in the PD Mgmt screen.
<Ctrl> <N>	Press <Ctrl> <N> to move to the next menu screen among the main menu screens: VD Mgmt , PD Mgmt , Ctrl Mgmt , and Foreign View .	Press <Ctrl> <N> on the VD Mgmt screen to move to the PD Mgmt screen.
<Ctrl> <P>	Press <Ctrl> <P> to move to the previous menu screen among the main menu screens: VD Mgmt , PD Mgmt , Ctrl Mgmt , and Foreign View .	Press <Ctrl> <P> on the PD Mgmt screen to move to the VD Mgmt screen.
<F1>	Press <F1> to access Help information. The Help screens display a glossary of topics you can use to access information about navigation, RAID levels, and general topics.	<F1>
<F2>	Press <F2> to access the context menu, which displays the list of options.	<F2>
<F5>	Press <F5> to refresh the information on the screen.	<F5>
<F11>	Switch between two controllers.	<F11>
<F12>	Press <F12> to display a list of controllers.	<F12>
Spacebar	Press the <spacebar> to select an item.	Press the <spacebar> to select or deselect a controller setting in the Ctrl Mgmt View .

Setting Up Virtual Disks


You can set up a disk group and create virtual disks using the procedures contained in this section. Each of the following procedures is explained individually in this section in detail.

To set up virtual disks:

1. Create the virtual disks and select the virtual disk options.
2. Designate hot spares (optional).


For more information, see the topic [Managing Dedicated Hot Spares](#).

3. Initialize the virtual disks.



 **NOTE:** When you use one physical disk group to create multiple virtual disks, all the virtual disks must be configured with the same RAID level.

When you define the virtual disks, you can set the following virtual disk parameters:

- RAID level
- Stripe element size
- Read policy
- Write policy
- Type of initialization
- Hot spare configuration

 **NOTE:** The default hard drive cache policy for a virtual disk composed with SAS hard drives is disabled and with SATA hard drives is enabled. The Virtual Disk parameter cannot be changed in the **BIOS Configuration Utility** (<Ctrl> <R>). Use Dell OpenManage Storage Management for the hard drive cache setting operation.


The following table shows the parameters that you can configure when defining virtual disks.


Parameter	Description
RAID Level	RAID Level specifies whether the virtual disk is RAID 0, 1, 5, 6, 10, 50, or 60. The number of disks, disk capacity, the requirements for fault tolerance, performance, and capacity should be considered when selecting the RAID level. For more information, see the topic Summary Of RAID Levels .
Stripe Element Size	Stripe Element Size specifies the size of the segments written to each physical disk in a RAID 0, 1, 5, 6, 10, 50, and 60 virtual disk. You can set the stripe element size to 64 KB 128 KB, 256 KB, 512 KB, or 1024 KB. The default and recommended stripe element size is 64 KB. A larger stripe element size provides better read performance if your system mostly does sequential reads.
Write Policy	Write Policy specifies the controller write policy. You can set the write policy to Write-Back or Write-Through. In Write-Back caching, the controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction.  NOTE: If a Battery Backup Unit (BBU) is present, the default cache setting is Write-Back . If no BBU is present, the default cache policy default setting is Write-Through .  NOTE: If Write-Back is enabled and the system is quickly turned off and then on, the controller may pause as the system flushes cache memory. Controllers that contain a battery backup default to Write-Back caching. In Write-Through caching, the controller sends a data transfer completion signal to the host when the disk subsystem has received all the data in a transaction.
Read Policy	Read-Ahead enables the read ahead feature for the virtual disk. You can set the parameter to Read-Ahead , No-Read-Ahead , or Adaptive . The default is Adaptive-Read-Ahead . Read-Ahead specifies that the controller uses Read-Ahead for the current virtual disk. Read-Ahead capability allows the controller to read sequentially ahead of requested data and store the additional data in cache memory, anticipating that the data is required soon.

Parameter	Description
	<p>No-Read-Ahead specifies that the controller does not use read ahead for the current virtual disk.</p> <p>Adaptive specifies that the controller begins using Read-Ahead if the two most recent disk accesses occurred in sequential sectors. If all read requests are random, the algorithm reverts to No-Read-Ahead; however, all requests are still evaluated for possible sequential operation.</p>

Virtual Disk Management


Creating Virtual Disks

 **NOTE:** Combining SAS and SATA disk drives within a virtual disk is not supported. Also, combining disk drives and SSDs within a virtual disk is not supported.


 **NOTE:** To create secured virtual disks, see the topic [Security Key And RAID Management](#).

Perform the following steps to create a virtual disk:

1. During host system bootup, press **<Ctrl><R>** after the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed. If there is more than one controller, the main menu screen is displayed. Select a controller, and press **<Enter>**. The **Virtual Disk Management** screen is displayed for the selected controller.
2. Use the arrow keys to highlight **Controller #** or **Disk Group #**.
3. Press **<F2>**.
The list of available actions is displayed.
4. Select **Create New VD** and press **<Enter>**.
The **Create New VD** screen is displayed. The cursor is on the **RAID Levels** option. When adding a virtual disk to a Disk Group, the **Add VD in Disk Group** screen is displayed. Skip to step 11 to change the basic settings of the virtual disk.
5. Press **<Enter>** to display the possible RAID levels, based on the physical disks available.
6. Press the down-arrow key to select a RAID level and press **<Enter>**.
7. When creating a spanned virtual disk (RAID 10, 50 or 60), enter the number of physical disks per span in the **PD per Span** field and press **<Enter>**.

 **NOTE:** Creating a 22 physical disk RAID 10 VD is possible by selecting RAID 10 and populating the **PD per Span** field with 22.

8. Press **<Tab>** to move the cursor to the list of physical disks.
9. Use the arrow key to highlight a physical disk and press the spacebar, **<Alt>**, or **<Enter>** to select the disk.
10. Select additional disks, if required.

 **NOTE:** (PERC H310) Only RAID Capable physical disks are eligible to be included in a Virtual Disk. To convert physical disks to RAID Capable, see the topic [Converting Physical Disk To RAID Capable For PERC H310](#).

11. Press **<Tab>** to move the cursor to the **Basic Settings** box.
12. Set the virtual disk size in the **VD Size** field.
The virtual disk size is displayed in GB format.
13. Press **<Tab>** to access the **VD Name** field, and type a virtual disk name.
14. Press **<Tab>** to move the cursor to **Advanced Settings**.

15. Press the spacebar to make the settings active so that you can change them.

An **X** is displayed beside **Advanced Settings**. The settings are the **stripe element size**, **read policy**, and **write policy**. You can also choose advanced options such as forcing the cache policy to **Write-Back**, initializing the virtual disk, and configuring a dedicated hot spare. The default parameters are displayed. You can accept the defaults or change them. To change the virtual disk parameters, see Virtual Disk Parameters and Descriptions in [Setting Up Virtual Disks](#).

Selecting Virtual Disk Parameters

To select the virtual disk parameters:

1. Press **<Tab>** to move the cursor to the parameters you want to change.
2. Press the down-arrow key to open the parameters and scroll down the list of settings.
3. To change the stripe element size, press **<Tab>** to highlight **Stripe Element Size**.
4. Press **<Enter>** to display the list of stripe element sizes (64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB). Press the down-arrow key to highlight the desired option and press **<Enter>**.
The default stripe element size is 64 KB.
5. If you need to change the read policy, press **<Tab>** to move the cursor to **Read Policy**.
6. Press **<Enter>** to display the options (**No Read Ahead**, **Read Ahead**, or **Adaptive Read Ahead**). Press the down-arrow key to highlight the desired option and press **<Enter>**.
7. If you need to change the write policy, press **<Tab>** to move the cursor to **Write Policy**.
8. Press **<Enter>** to display the options (**Write-Through**, **Write Back**). Press the down-arrow key to highlight an option and press **<Enter>**.
9. Press **<Tab>** to move the cursor to **Force WB with no battery** and press **<Enter>**. If you chose **Write-Through** as the write policy, then the option is not available.



CAUTION: Do not initialize virtual disks when attempting to recreate an existing configuration.

10. Press **<Tab>** to move the cursor to **Initialize** and press **<Enter>**.



NOTE: The initialization performed at the stage is fast initialization.

11. Press **<Tab>** to move the cursor to **Configure Hot Spare** and press **<Enter>**.



NOTE: The hot spare created at the stage is a dedicated hot spare.

12. If you have chosen to create hot spares in the earlier steps, a pop-up window is displayed where disks with appropriate sizes are displayed. Press the **<spacebar>** to select the disk size.
13. After you select the disk size, click **OK** to finalize the selection or click **Cancel** to forfeit the selection.
14. Select **OK** to accept the settings and press **<Enter>** to exit the window or select **Cancel** and press **<Enter>** to exit if you do not want to change any virtual disk parameters.

Converting Physical Disk To RAID Capable For PERC H310



NOTE: By default, all physical drives are discovered as Non-RAID drives.

Conversion of a physical disk to RAID capable allows that disk to be used within a Virtual Disk or as a Hot Spare.

1. Press **<Ctrl> <N>** to access the **PD Mgmt** screen.
A list of physical disks is displayed. The status of the each disk is displayed under **State**.
2. Press the down-arrow key to highlight a physical disk.
3. Press **<F2>** to display the menu of available actions.
4. Press the down-arrow key to highlight **Convert to RAID Capable**.

5. Press **<Enter>**.

Converting Physical Disk To Non-RAID For PERC H310

Conversion of a Physical Disk to Non-RAID allows direct operating system access to the drive. Non-RAID disks cannot be used within a Virtual Disk nor can they be used as Hot Spares.


1. Press **<Ctrl> <N>** to access the **PD Mgmt** screen.
A list of physical disks is displayed. The status of the each disk is displayed under **State**.
2. Press the down-arrow key to highlight a physical disk.
3. Press **<F2>** to display the menu of available actions.
4. Press the down-arrow key to highlight **Convert to Non-RAID**.
5. Press **<Enter>**.

Initializing Virtual Disks

 **CAUTION:** A full initialization permanently destroys any existing data on that virtual disk.

Perform the following steps to initialize virtual disks:

1. On the **VD Mgmt screen**, select **Virtual Disk #** and press **<F2>** to display the menu of available actions.
2. Select **Initialization** and press the right-arrow key to display the **Initialization** submenu options.
3. Select **Start Init.** to begin a regular initialization or select **Fast Init.** to begin a fast initialization.
A pop-up window is displayed indicating that the virtual disk has been initialized.
4. Repeat the procedures from step 1 to step 3 to configure another virtual disk.

 **NOTE:** The PERC H310 card supports up to 16 virtual disks per controller, and the PERC H710, H710P, and H810 cards support up to 64 virtual disks per controller. The currently configured virtual disks display on the screen.

Checking Data Consistency

Select the **Consistency Check (CC)** option in the configuration utility to verify the redundancy data in virtual disks that use RAID levels 1, 5, 6, 10, 50, and 60. (RAID 0 does not provide data redundancy.)

If you attempt to run a **Consistency Check** on a virtual disk that has not been initialized, the following error message is displayed:

The virtual disk has not been initialized. Running a consistency check may result in inconsistent message in the log. Are you sure you want to continue?

You can select **Yes** or **No**. If you select **Yes**, the CC operation continues. If you select **No**, the operation ends.

Running A Data Consistency Check

To perform a consistency check:

1. Press **<Ctrl> <N>** to access the **VD Mgmt** menu screen.
2. Press the down-arrow key to highlight **Virtual Disk #**.
3. Press **<F2>** to display the menu of available actions.
4. Press the down-arrow key to select **Consistency Check**.
5. Press the right-arrow key to display the available actions (**Start, Stop**).
6. Select **Start** and press **<Enter>** to run a **Consistency Check**.


The **Consistency Check** runs and checks the redundancy data in the virtual disks.


7. After you start the **Consistency Check**, press <Esc> to display the previous menu if needed.

Importing Or Clearing Foreign Configurations Using The VD Mgmt Menu

When a foreign configuration exists, the BIOS screen displays the message `Foreign configuration(s) found on adapter`. In addition, a foreign configuration is displayed on the right side of the **Ctrl Mgmt** screen.


You can use the **VD Mgmt** menu to import the existing configuration to the RAID controller or clear the existing configuration. In addition, you can view the foreign configuration from the **Foreign View** tab without importing the configuration.

 **NOTE:** The controller does not allow an import of configurations that results in more than 64 virtual disks.

 **NOTE:** To import a secured foreign configuration, see the topic [Security Key And RAID Management](#).

Perform the following steps to import or clear foreign configurations:


1. During bootup, press <Ctrl> <R> when prompted by the BIOS screen.
The **VD Mgmt** screen is displayed by default.
2. On the **VD Mgmt** screen, highlight the **Controller #**.
3. Press <F2> to display the available actions.
4. Navigate to the **Foreign Config** option and press the right arrow key to display the available actions: **Import** and **Clear**.

 **NOTE:** Ensure that your virtual disk has all the physical disks by verifying that there are no physical disks marked as **Missing** in the foreign view page and that all the disks appear as expected before importing them.

5. Select **Import** to import the foreign configuration or **Clear** to delete the foreign configuration and then press <Enter>.


If you import the configuration, the **VD Mgmt** screen displays detailed configuration information. It includes information about the disk groups, virtual disks, physical disks, space allocation, and hot spares.

Importing Or Clearing Foreign Configurations Using The Foreign Configuration View Screen

 **NOTE:** To import a secured foreign configuration, see the topic [Security Key And RAID Management](#).

If one or more physical disks are removed from a configuration, the configuration on those disks is considered a foreign configuration by the RAID controller.

You can use the **Foreign Configuration View** screen to view information about the foreign configuration, such as disk groups, virtual disks, physical disks, space allocation, and hot spares. The foreign configuration data is displayed in the same format as configurations on the **VD Mgmt** screen. You can use the **VD Mgmt** screen to view the foreign configuration before importing. After you view the foreign configuration, you can either clear or import to the RAID controller.

 **NOTE:** Before you import the foreign configuration, review the configuration on the screen to ensure that it is the end result that you require.

You can use the **Foreign Configuration View** screen to manage foreign configurations in the following cases:

- All the physical disks in a configuration are removed and re-inserted.
- Some of the physical disks in a configuration are removed and re-inserted.
- All the physical disks in a virtual disk are removed, but at different times, and then re-inserted.

- The physical disks in a non-redundant virtual disk are removed.


The following constraints apply to the physical disks that are considered for import:

- The disk state of a physical disk can change from the time the foreign configuration is scanned to when the actual import occurs. The foreign import occurs only on disks that are in the **Unconfigured Good** state.
- Disks in the failed or offline state cannot be imported.
- The firmware does not allow you to import more than eight foreign configurations.

Perform the following procedures on the **Foreign Configuration View** screen to manage foreign configurations in each specific case:


1. If all or some of the physical disks in a configuration are removed and reinserted, the controller considers the disks to have foreign configurations. Perform the following steps:


- a) Select **Foreign Configuration View** to display the foreign configuration information on the **Foreign Configuration View** screen.
- b) Press <F2> to display the options (**Import**, **Clear**).

 **NOTE:** You must have all the disks in the system before you perform the import operation.

- c) Select **Import** to import the foreign configuration to the controller or select **Clear** to delete the foreign configuration(s) from the re-inserted disk(s).


In the **Preview Configuration Data** window, the status of a physical disk that needs to be rebuilt is displayed as **Rebuild**.

 **NOTE:** When you import a foreign configuration, the dedicated hot spares in the configuration are imported as dedicated hot spares on two conditions — the associated virtual disk is already present or the associated virtual disk is also imported along with the configuration.

 **NOTE:** Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual disks. For more information about checking data consistency, see the topic [Checking Data Consistency](#).


2. If all the physical disks in a virtual disk are removed at different times and re-inserted, the controller considers the disks to have foreign configurations. Perform the following steps:

- a) Select **Foreign Configuration View** to display the complete virtual disk, across different foreign configurations and allow foreign configurations to be imported.
- b) Press <F2> to display the options **Import** and **Clear**.

 **NOTE:** You must have all the drives in the system before you perform the import operation.

- c) Select **Import** to merge the foreign configurations with the existing configuration on the controller or **Clear** to delete the foreign configuration(s) from the re-inserted disk(s).

If you select **Import**, all drives that were pulled before the virtual disk became offline are imported, and then automatically rebuilt.

 **NOTE:** Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual disks. For more information about checking data consistency, see the topic [Checking Data Consistency](#).

3. If the physical disks in a non-redundant virtual disk are removed, the controller considers the disks to have foreign configurations. Perform the following steps:

- a) Select **Foreign Configuration View** to display the complete foreign configuration information.
- b) Press <F2> to display the options **Import** and **Clear**.
- c) Select **Import** to import the foreign configurations to the virtual disk or **Clear** to delete the foreign configuration(s) from the re-inserted disk(s).

No rebuilds occur after the import operation because there is no redundant data to rebuild the disks with.

Break Mirror

A Break Mirror operation is an operation that can be performed only on RAID1 arrays. It provides a way to 'split' the mirror and spin-down one of the hard disks, which can then be imported into the configuration of a different PERC H310, H710, H710P, or H810 controller. This can be a useful way to:

- Create a disk image that can be imported and booted on a different system.
- Assist in software or configuration testing, where half of the mirror can be removed to ensure configuration consistency.



NOTE: The **Break Mirror** operation is not available within the booted operating system environment. It is only available in the BIOS (<Ctrl> <R>) utility and the UEFI RAID configuration utility.

Break Mirror In BIOS Configuration Utility

The RAID1 VD that you are to split must be in an optimal state, and it cannot be running any other background tasks.

To use the **Break Mirror** feature from the **BIOS Configuration Utility** (<Ctrl> <R>), perform the following steps:

1. During boot up, press <Ctrl> <R> when prompted by the BIOS screen. If there is more than one controller in the system, choose the controller.

The **VD Mgmt** screen is displayed by default.

2. Highlight the appropriate **Disk Group**.
3. Press <F2> to display the available actions.
4. Choose **Break Mirror** and press <Enter>.

A dialog is displayed describing the operation and asking if you are sure you want to perform the operation.

5. Select **Yes** to continue.

The exported disk (exporting configuration) spins down and the hard drive status LED starts flashing to identify which drive must be pulled.

The remaining disk (exporting configuration) is in a degraded state until the missing member is replaced or rebuilt.

If there are any Global or Dedicated Hot Spare assigned which fulfills the rebuild requirements for the degraded RAID1 disk, a rebuild starts automatically. If there is no Hot Spare assigned, then a Hot Spare that meets all the requirements for the VD must be assigned before the rebuild starts.



NOTE: Importing a Broken Mirror is the same as importing a Foreign configuration. See the topic [Importing Or Clearing Foreign Configurations Using The VD Mgmt Menu](#). The imported VD is in a degraded state until the missing member is rebuilt.

Managing Preserved Cache

If a virtual disk becomes offline or is deleted because of missing physical disks, the controller preserves the dirty cache from the virtual disk. The preserved dirty cache, known as pinned cache, is preserved until you import the virtual disk or discard the cache.



NOTE: Certain operations, such as creating a new virtual disk, cannot be performed if preserved cache exists. You have to enter the **BIOS Configuration Utility** (<Ctrl> <R>) to resolve the situation before you boot to the operating system. Messages are displayed notifying you that you must enter the **BIOS Configuration Utility** (<Ctrl> <R>) to discard the preserved cache or import the virtual disks with the preserved cache.



CAUTION: If there are any foreign configurations, it is strongly advised that you import the foreign configuration before you discard the preserved cache. Otherwise, you might lose data that belongs with the foreign configuration.

Perform the following steps to select whether to import the virtual disk or discard the preserved cache:




1. On the **VD Mgmt** screen, click on a controller icon.
2. Press <F2> to display the available actions.
3. Select **Manage Preserved Cache**.
A message is displayed advising you to import the foreign configuration before you discard the preserved cache to avoid losing data belonging to the foreign configuration. Confirm whether you want to continue. The **Manage Preserved Cache** screen displays the affected virtual disks.
4. You can choose to discard the cache on the **Manage Preserved Cache** screen. If you press **Cancel**, the process is cancelled and the **Preserved Cache Retained** dialog box is displayed.
If you choose to discard the cache, you are prompted to confirm your choice. If you choose to retain the cache, a message is displayed to notify you that you cannot perform certain operations while the cache exists. Click **OK** to continue.

Managing Dedicated Hot Spares

A dedicated hot spare automatically replaces a failed physical disk only in the selected disk group which the hot spare is part of.

A dedicated hot spare is used before a global hot spare is used. You can create dedicated hot spares or delete them on the **VD Mgmt** screen.

Perform the following steps to create or delete dedicated hot spares:

1. On the **VD Mgmt** screen, select **Disk Group #** and press <F2> to display the menu of available actions.
The available menu options are displayed.
2. Select **Manage Ded. HS** and press <Enter>.
A screen displays a list of the current dedicated hot spares with an **X** beside them and the physical disks that are available to create dedicated hot spares.
 **NOTE:** The utility allows only disks of the same disk technology and of equal or greater size to be selected as dedicated hot spare.
3. Use the following instructions to create or delete a dedicated hot spare:
 - Creating a dedicated hot spare
 1. Press the down-arrow key to highlight an available physical disk.
 2. Press the spacebar to select the disk.
 3. Repeat step a to step b for each dedicated hot spare that you want to create. **NOTE:** An **X** is displayed beside the selected physical disk(s).
 - Deleting a dedicated hot spare
 1. Use the down-arrow key to highlight a current hot spare.
 2. Press the spacebar to de-select the disk.
 3. Repeat step a and step b for each dedicated hot spare that you want to delete.
4. Press <Enter> to confirm the changes.
The **VD Mgmt** screen displays the updated list of hot spares under the **Hot spares** heading.
 **NOTE:** If a global hot spare or dedicated hot spare is removed, reinserted and imported, it regains its status as a hot spare. A dedicated hot spare becomes a global hot spare when the disk group it was assigned to protect is no longer present during import.

Deleting Virtual Disks



NOTE: You cannot delete a virtual disk during an initialization.



NOTE: Warning messages appear stating the effect of deleting a virtual disk. Click **OK** twice to complete the virtual disk deletion.

To delete virtual disks, perform the following steps in the **BIOS Configuration Utility** (<Ctrl> <R>):

1. Press <Ctrl> <N> to access the **VD Mgmt** screen.
2. Use the arrow keys to move the cursor to the **Virtual Disks** heading.
3. Press <F2>.
The action menu is displayed.
4. Select **Delete VD** and press <Enter>.
5. If there are multiple virtual disks in a **Disk Group**, select **Total Free Capacity** for the Disk Group in the **VD Mgmt** screen.
The total amount of free space available in the **Disk Group** is displayed.

Deleting Disk Groups

You can delete disk groups using the **BIOS Configuration Utility** (<Ctrl> <R>). When you delete a disk group, the utility also removes the virtual disks in that disk group.

To delete disk groups, perform the following steps in the **BIOS Configuration Utility** (<Ctrl> <R>):

1. Press <Ctrl> <N> to access the **VD Mgmt** screen.
2. Use the arrow keys to move the cursor to the **Virtual Disks** heading.
3. Press <F2>.
The action menu is displayed.
4. Select **Delete Disk Group** and press <Enter>.
The disk group is deleted.



NOTE: When you delete a disk group, the remaining disk groups with higher numbers are automatically renumbered. For example, if you delete disk group #2, then disk group #3 is automatically renumbered as disk group #2.

Clearing The Configuration

You can delete all virtual disks on the RAID controller by performing the operation.

To clear the configuration, perform the following steps in the **BIOS Configuration Utility** (<Ctrl> <R>):

1. Press <Ctrl> <N> to access the **VD Mgmt** screen.
2. Use the arrow keys to move the cursor to the **Controller** heading.
3. Press <F2>.
The action menu is displayed.
4. Select **Clear Config**.
A pop-up window is displayed prompting for confirmation to delete all virtual disks.
5. Select **NO** to delete the virtual disks or **YES** to retain the [Consistency Checks](#) existing configuration.

BIOS Configuration Utility Menu Options

The first menu that is displayed when you access the **BIOS Configuration Utility** <Ctrl> <R> is the main menu screen. It lists the controller, controller number, and other information, such as the slot number. On the screen, you can use the arrow keys to select the RAID controller you want to configure. Press <Enter> to access the controller

This section describes the options for the **BIOS Configuration Utility** (<Ctrl> <R>) for each of the major menus:

- **Virtual Disk Management (VD Mgmt)** menu
- **Physical Disk Management (PD Mgmt)** menu
- **Controller Management (Ctrl Mgmt)** menu
- **Foreign Configuration View (Foreign View)** menu

Most menus consist of two panels:

- A left panel with the menu options
- A right panel with details of the items selected in the left panel

The following sections describe the menu and submenu options for each of the major menus:

Virtual Disk Management (VD Mgmt)

The Virtual Disk Management screen, VD Mgmt, is the first screen that is displayed when you access a RAID controller from the main menu screen on the **BIOS Configuration Utility** (<Ctrl> <R>). The left panel displays the menus for the virtual disk management as given below:

Controller—Consists of the sub-menu item **Disk Group** which again has the following sub-menu items:

- **Virtual Disks**
- **Physical Disks**
- **Total Free Capacity** (virtual disk size and free space you can use to create a virtual disk)
- **Hot Spares** (global and dedicated)

The right panel displays detailed information for the selected controllers, disk groups, virtual disks, physical disks, total free capacity, and hot spares, as shown in the following table:

Table 5. Information on the Virtual Disk Management Screen

Menu Item Selected in Left Panel	Information Displayed in Right Panel
Controller	Controller Properties: <ul style="list-style-type: none">• Number of disk groups (DG)• Number of virtual disks (VD)• Number of physical disks (PD)
Disk Group #	Disk Group # Properties: <ul style="list-style-type: none">• Number of virtual disks (VD)• Number of physical disks (PD)• Space available on the physical disks• Number of free segments• Number of dedicated hot spares

Menu Item Selected in Left Panel	Information Displayed in Right Panel
	<ul style="list-style-type: none"> • Security property of the Disk Group
Virtual Disks	<p>Disk Group # Properties:</p> <ul style="list-style-type: none"> • Number of virtual disks (VD) • Number of physical disks (PD) • Space available in the virtual disk • Number of free segments • Number of dedicated hot spares
Virtual Disk #	<p>Virtual Disk # Properties:</p> <ul style="list-style-type: none"> • RAID level (0, 1, 5, 6, 10, 50, or 60) • RAID status of the virtual disk (Failed, Degraded, Partially Degraded, or Optimal) • Operation currently in progress <p>Disk Group # Properties:</p> <ul style="list-style-type: none"> • Number of virtual disks (VD) • Number of physical disks (PD) • Space available on the physical disks • Number of free segments • Number of dedicated hot spares
Physical Disks	<p>Disk Group # Properties:</p> <ul style="list-style-type: none"> • Number of virtual disks (VD) • Number of physical disks (PD) • Space available on the physical disks • Number of free segments • Number of dedicated hot spares
Physical Disk #	<p>Physical Disk Properties:</p> <ul style="list-style-type: none"> • Vendor name • Physical disk state • Enclosure Position • Slot Position <p>Disk Group # Properties:</p> <ul style="list-style-type: none"> • Number of virtual disks (VD) • Number of physical disks (PD) • Space available on the physical disks • Number of free segments • Number of dedicated hot spares
Total Free Capacity	<p>Disk Group # Properties:</p> <ul style="list-style-type: none"> • Number of virtual disks (VD)

Menu Item Selected in Left Panel	Information Displayed in Right Panel
	<ul style="list-style-type: none"> • Number of physical disks (PD) • Space available on the physical disks • Number of free segments • Number of dedicated hot spares
Hot Spares	<p>Physical Disk Properties:</p> <ul style="list-style-type: none"> • Vendor name • Physical disk state • Enclosure Position • Slot Position <p>Disk Group # Properties:</p> <ul style="list-style-type: none"> • Number of virtual disks (VD) • Number of physical disks (PD) • Space available on the physical disks • Number of free segments • Number of dedicated hot spares

Virtual Disk Actions

The following table describes the actions you can perform on virtual disks. For procedures you can use to perform the actions, see [Virtual Disk Management](#).

Table 6. Virtual Disk Actions

Action	Description
Create a new virtual disk	Creates a new virtual disk from one or more physical disks. You can configure hot spares when you create a virtual disk.
Manage dedicated hot spares	Creates or deletes a hot spare that you can dedicate to a single redundant virtual disk.
Initialize a virtual disk	Initializes the selected virtual disk. You must initialize every virtual disk that is configured. You can perform a fast initialization or a full initialization.
Check data consistency on a virtual disk	Verifies the correctness of the redundancy data in the selected virtual disk. The option is available only if RAID level 1, 5, 6, 10, 50, or 60 is used. The PERC H310, H710, H710P, or H810 cards automatically correct any differences found in the data.
Display or update virtual disk parameters	Displays the properties of the selected virtual disk. You can modify the cache write policy and read policy from the menu.
Manage preserved cache	Preserves the dirty cache from a virtual disk if it becomes offline or is deleted. The dirty cache is preserved until you import the virtual disk or discard the cache.
Delete a virtual disk	Deletes the virtual disk and frees up disk space to create another virtual disk.
Delete a disk group	Deletes a disk group, which is a collection of disks from one or more disk subsystems controlled by management software.

Physical Disk Management (PD Mgmt)

The **Physical Disk Management screen (PD Mgmt)** displays physical disk information and action menus. The screen displays physical disk IDs, vendor names, disk size, type, state, and disk group (DG). You can sort the list of physical disks based on the headings. You can perform several actions on the physical disks, including the following:

- Rebuilding physical disks
- Performing the Replace Member operation
- Setting the LED to blink
- Making a disk online or offline (unaffiliated with a disk group)
- Creating global hot spares
- Removing dedicated hot spares or global hot spares

The **PD Mgmt** screen also displays several physical disk properties as shown in the following table.

Table 7. Information on the Physical Disk Management Screen

Information Displayed in Left Panel	Supported Information Displayed in Right Panel
Physical Disk: <ul style="list-style-type: none">• Disk ID• Protocol type• Capacity (GB)• Physical Disk State• Disk Group• Vendor	<ul style="list-style-type: none">• Security Property of Physical Disk• Encryption Capable• Product ID• Firmware Revision• Disk Write Cache• S.M.A.R.T state• Physical Disk operation• Max Device Link Rate• Negotiated Link Rate• Dell Certified Disk

Physical Disk Actions

The following table describes the actions you can perform on physical disks. For procedures that can be used to perform the actions, see the topic [Physical Disk Management \(PD Mgmt\)](#).

Table 8. Physical Disk Actions

Action	Description
Rebuild	Regenerates all data to a replacement disk in a redundant virtual disk (RAID level 1, 5, 6, 10, 50, or 60) after a disk failure. A disk rebuild normally occurs without interrupting normal operations on the affected virtual disk.
Replace Member	Replaces the disk in the virtual disk with another disk that can be selected.
LED Blinking	Indicates when physical disks are being used to create a virtual disk. You can choose to start or stop the LED blinking.
Force Online	Changes the state of the selected physical disk to online.
Force Offline	Changes the state of the selected physical disk so that it is no longer part of a virtual disk.
Make Global HS	Designates the selected physical disk as a global hot spare. A global hot spare is part of a pool for all virtual disks controlled by the controller.□

Action	Description
Remove HS	Removes a dedicated hot spare from its disk group or a global hot spare from the global pool of hot spares.

Rebuild

Select **Rebuild** to rebuild one or more failed physical disks. For information on performing a physical disk rebuild, see the topic [Performing A Manual Rebuild Of An Individual Physical Disk](#).

Several of the controller configuration settings and the virtual disk settings affect the actual rate of rebuild. The factors include the rebuild rate setting, virtual disk stripe size, virtual disk read policy, virtual disk write policy, and the amount of workload placed on the storage subsystem. For information on getting the best rebuild performance from your RAID controller, see the documentation at dell.com/support/manuals.

The listed rates in the following table were taken during single disk failure with no I/O present on a PERC H810 card connected to a single PowerVault MD1220 enclosure. Rates vary depending on type, speed and number of hard drives present in array; as well as which controller model and enclosure configuration are being used.

Table 9. Estimated Rebuild Rates

RAID Level	Number of <input type="checkbox"/> Hard Drives	7.2K RPM 6 Gb/s <input type="checkbox"/> SAS Hard Drive	15K RPM 6 Gb/s <input type="checkbox"/> SAS Hard Drive
RAID 1	2	320 GB/hour	500 GB/hour
RAID 5	6	310 GB/hour	480 GB/hour
RAID 10	6	320 GB/hour	500 GB/hour
RAID 5	24	160 GB/hour	240 GB/hour
RAID 10	24	380 GB/hour	500 GB/hour

Controller Management (Ctrl Mgmt)


The **Controller Management** screen (**Ctrl Mgmt**) displays the product name, package, firmware version, BIOS version, boot block version, controller ID, security capability, and security key presence. Use the screen to perform actions on the controller and BIOS. You can perform functions such as enable or disable the controller BIOS, enable or disable the BIOS during bootup in the event of BIOS errors, and enable or disable the option to **Auto Import**. In addition, you can select a virtual disk from which to boot, and select default settings.

Controller Management Actions

The following table describes the actions you can perform on the **Ctrl Mgmt** screen.

Table 10. Controller Management Options

Option	Description
Enable Controller BIOS	<p>Select the option to enable the controller BIOS. If the boot device is on the RAID controller, the BIOS must be enabled.</p> <p>Disable the BIOS to use other boot devices. <input type="checkbox"/></p> <p>In a multiple controller environment, you can enable BIOS on multiple controllers. However, if you want to boot from a specific controller, then enable the BIOS on that</p>

Option	Description
Enable BIOS Stop On Error	controller and disable it on the other controllers. The system can then boot from the BIOS-enabled controller. Select the option to stop the system BIOS during bootup if there are BIOS errors. The option enables you to enter the configuration utility to resolve the problem.
Select Bootable Device	Select the option to specify a virtual disk as the boot disk on the controller.  NOTE: For PERC H310 the option may be a Virtual Disk or <input type="checkbox"/> Non-RAID drive. The option is displayed if you have built virtual disks.
Enable Auto Import	Attempts to import every online foreign configuration during boot without having the need to access the BIOS Configuration Utility (<Ctrl> <R>).
Factory Default	Select the option to restore the default settings for the options in the Settings box.

Foreign Configuration View

When a foreign configuration is present, you can select **Foreign Configuration View** to display the configuration. The screen shows the foreign configuration as it would be if you import it. You can preview the foreign configuration before you decide whether to import it or clear it.

In some cases, a foreign configuration cannot be imported. If a physical disk in a virtual disk is rebuilding, the physical disk's state is set to **Rebuild**. No virtual disk target ID displays for virtual disks that cannot be imported.

The section Importing or Clearing Foreign Configurations Using the Foreign Configuration View Screen contains the procedures you can use to manage the foreign configurations.

 **NOTE:** The BIOS Configuration Utility (<Ctrl> <R>) reports error codes for failed imports of foreign configurations.

Physical Disk Management

Setting LED Blinking

The LED blinking option indicates when physical disks are being used to create a virtual disk. You can choose to start or stop the LED blinking.

Perform the following steps to start or stop the option:


1. Press <Ctrl> <N> to access the **PD Mgmt** screen.
A list of physical disks is displayed. The status of the each disk is displayed under the heading **State**.
2. Press the down-arrow key to highlight a physical disk.
3. Press <F2> to display the menu of available actions.
4. Press the down-arrow key to highlight **LED Blinking**.
5. Press the right-arrow key to display the available actions, **Start** and **Stop**.
6. Select **Start** to begin LED blinking or **Stop** to end LED blinking.

Creating Global Hot Spares

You can use a global hot spare to replace a failed physical disk in any redundant array as long as the capacity of the global hot spare is equal to or larger than the coerced capacity of the failed physical disk.

Perform the following steps to create global hot spares:

1. Press the down-arrow key to highlight a physical disk to change to a global hot spare.
2. Press <F2> to display the menu of available actions.
3. Press the down-arrow key to highlight **Make Global HS** and press <Enter>. The physical disk is changed to a global hot spare. The status of the physical disk as a global hot spare is displayed under the heading **State**.


 **NOTE:** To replace a failed physical disk global hot spares must use the same disk technology and must be equal or greater in size.

4. Select additional physical disks if desired and follow the previous steps to change them to global hot spares.

Removing Global Or Dedicated Hot Spares

You can remove one global or dedicated hot spare at a time on the **PD Mgmt** screen. Perform the following steps to remove a global hot spare or dedicated hot spare:

1. Press <Ctrl> <N> to access the **PD Mgmt** screen. A list of physical disks is displayed. The status of the each disk is displayed under the heading **State**.
2. Press the down-arrow key to highlight a physical disk that is a hot spare.
3. Press <F2> to display the menu of available actions.
4. Press the down-arrow key to select **Remove Hot Spare** from the list of actions and press <Enter>. The physical disk is changed to the **Ready** state. The status of the physical disk is displayed under the heading **State**.

 **NOTE:** Try to use physical disks of the same capacity in a specific virtual disk. If you use physical disks with different capacities in a virtual disk, all physical disks in the virtual disk are treated as if they have the capacity of the smallest physical disk.

5. Select additional hot spares if desired and follow step 1 to step 4 to remove them.

Replacing An Online Physical Disk

In addition to the automatic **Replace Member** operation, you can manually replace any physical disk that is part of a virtual disk using the **Replace Member** functionality.

Perform the following steps to replace a physical disk:

1. In the **Virtual Disk Management** window, select **Virtual Disk #** and press the down-arrow key until **Physical Disks** is highlighted.
2. Press the right-arrow key to expand the list of physical disks that are members of the virtual disk.
3. Press the down-arrow key and highlight the desired physical disk you want to replace. Press <F2> to expand the list of allowed operations on the disk.
4. Select **Replace** and then **Start**.
5. Press the down-arrow to highlight a replacement disk and then press the spacebar to select the disk.
6. Select **OK** to start the replacement.



NOTE: The replacement disk must be a hot spare or an unconfigured disk without a foreign configuration. It must have the same or greater capacity and should be of the same type as the disk it is replacing.

Restrictions and Limitations

The following restrictions and limitations apply to the **Replace Member** operation:

- The **Replace Member** functions are restricted to one per array for RAID 0, RAID 1, and RAID 5, and two per array for RAID 6.
- The **Replace Member** function and rebuild cannot run simultaneously on a RAID 6 virtual disk. The rebuild operation has a higher priority, and the **Replace Member** operation is aborted if a rebuild begins.
- The **Replace Member** function and rebuild cannot run simultaneously on a RAID 6 virtual disk. The rebuild operation has a higher priority, and the **Replace Member** operation is aborted if a rebuild begins.

Stopping Background Initialization

Background initialization (BGI) is the automated operation in which parity is created and written. BGI does not run on RAID 0 virtual disks. Under certain conditions, the **BIOS Configuration Utility (<Ctrl> <R>)** displays a message if you want to stop BGI in progress. An alert message is displayed if BGI is in progress and you start any of the following actions:

- A Full Initialization on the virtual disk
- A Fast Initialization on the virtual disk
- A Consistency Check on the virtual disk

The following alert message is displayed:

```
The virtual disk is undergoing a background initialization process. Would you
like to stop the operation and proceed with the <full initialization/quick
initialization/consistency check> instead?
```

Click **Yes** to stop the BGI and start the requested operation or **No** to allow BGI to continue.

Performing A Manual Rebuild Of An Individual Physical Disk



CAUTION: If a physical disk is a member of a disk group that contains multiple virtual disks and one of the virtual disks is deleted when a rebuild operation is in progress, the rebuild operation stops. You can then resume the rebuild operation manually using a storage management application. To avoid interruption, ensure that none of the virtual disks are deleted until the rebuild is complete.

Use the following procedures to rebuild one failed physical disk manually.

1. Press **<Ctrl> <N>** to access the **PD Mgmt** screen.
A list of physical disks is displayed. The status of each disk is displayed under the heading **State**.
2. Press the down-arrow key to highlight a physical disk that has a failed state.
3. Press **<F2>** to display a menu of available actions.
The **Rebuild** option is highlighted at the top of the menu.
4. Press the right-arrow key to display the rebuild options and select **Start**.
5. After you start the rebuild, press **<Esc>** to display the previous menu.



NOTE: You can also use the **VD Mgmt** screen to perform a manual rebuild. Use the arrow key to highlight a physical disk, and press **<F2>**. In the menu that is displayed, select the **Rebuild** option.

Controller Management

Enabling Boot Support



NOTE: See your system documentation to ensure the proper boot order is selected in the system BIOS.

In a multiple controller environment, you can enable BIOS on multiple controllers. However, if you want to boot from a specific controller, enable the BIOS on that controller and disable it on the other controllers. The system can then boot from the BIOS-enabled controller. Perform the following steps to enable the controller BIOS:

1. Press **<Ctrl> <N>** to access **Ctrl Mgmt** menu screen.
2. Press **<Tab>** to move the cursor to **Enable Controller BIOS** in the **Settings** box.
3. Press the spacebar to select **Enable Controller BIOS**.
An **X** is displayed beside **Enable Controller BIOS**.
4. Press **<Tab>** to move the cursor to the **Apply** button, and then press **<Enter>** to apply the selection.
The controller BIOS is enabled.
To disable the controller BIOS, use the spacebar to de-select the **Enable Controller BIOS** control, and then select **Apply** and press **<Enter>**.

Enabling Boot Support For A BIOS-Enabled Controller

1. Press **<Ctrl> <N>** to access the **Ctrl Mgmt** menu screen.
2. Press **<Tab>** to move the cursor to the **Select Boot Device** in the **Settings** box.
3. Press the down-arrow key to display a list of virtual disks and Non-RAID physical disks (PERC H310 only).
4. Use the down-arrow key to highlight a virtual disk or Non-RAID physical disk (PERC H310 only).
5. Press **<Enter>** to select the disk.
6. Press **<Tab>** to move the cursor to the **Apply** button, and then press **<Enter>** to apply the selection.
Boot support is enabled for the selected controller.

Enabling BIOS Stop On Error

The option **BIOS Stop on Error** is used to stop the system from booting if there are BIOS errors. Perform the following steps to enable **BIOS Stop on Error**.

1. Press **<Ctrl> <N>** to access the **Ctrl Mgmt** menu screen.
2. Press **<Tab>** to move the cursor to **Enable BIOS Stop on Error** in the **Settings** box.
3. Press the spacebar to select **Enable BIOS Stop on Error**.
An **X** is displayed beside **Enable BIOS Stop on Error**.
4. Press **<Tab>** to move the cursor to the **Apply** button, and then press **<Enter>** to apply the selection.
The controller BIOS is enabled. To disable **Enable BIOS Stop on Error**, use the spacebar to de-select **Enable BIOS Stop on Error**, then select **Apply** and press **<Enter>**.

Enabling Auto Import

If there is a native configuration present on the controller, the option **Enable Auto Import** automatically imports every online foreign configuration during boot without having the need to access the **BIOS Configuration Utility** (**<Ctrl> <R>**).



NOTE: The controller automatically imports every optimal and degraded foreign configuration without enabling the feature if there is no native configuration on the controller.

To enable Auto Import:

1. Press **<Ctrl> <N>** to access the **Ctrl Mgmt** menu screen.
2. Press **<Tab>** to move the cursor to **Enable Auto Import** in the **Settings** box.
3. Press the spacebar to select **Enable Auto Import**.
An **X** is displayed beside **Enable Auto Import**.
4. Press **<Tab>** to move the cursor to the **Apply** button, and then press **<Enter>** to apply selection.
The **Auto Import** is enabled.

Disabling Auto Import

To disable Auto Import:

1. Use the spacebar to de-select **Enable Auto Import**.
2. Select **Apply** and press **<Enter>**.
The **Auto Import** is disabled.

Restoring Factory Default Settings

You can use the **Ctrl Mgmt** menu screen to restore the default settings for the options in the **Settings** box. The settings are **Enable Controller BIOS**, **Enable BIOS Stop on Error**, and **Enable Auto Import**.

Perform the following steps to restore default settings:

1. Press **<Ctrl> <N>**.
2. Press **<Tab>** to move the cursor to the **Settings** box.
3. Use the spacebar to de-select the settings for the options in the **Settings** box.
4. Press **<Tab>** to move the cursor to the **Factory Default** box, and press **<Alt>**, **<Enter>**, or the spacebar.
A dialog box is displayed for you to confirm your choice.
5. Select **Ok** and press **<Enter>**.
The defaults are automatically selected for the controller settings and are displayed in **Settings**.

UEFI RAID Configuration Utility

The **UEFI RAID Configuration Utility** is a storage management application integrated into the System BIOS that can be used to configure and manage RAID disk groups, virtual disks and physical disks. This utility is independent of the operating system.



NOTE: Use the **UEFI RAID Configuration Utility** for initial setup and disaster recovery. Certain advanced features are also offered in the utility.

Entering The UEFI RAID Configuration Utility


To enter the **UEFI RAID Configuration** utility:


1. Turn on the system.
2. During startup, press **<F2>** to enter the **System Setup**.


3. From **System Setup Main Menu**, enter **Device Settings**.

Device Settings lists all the RAID controllers in the system.

Use the arrow keys or mouse to select the RAID controller you want to configure/manage and press **<Enter>** or click on the mouse to access the management menus for the controller.

 **NOTE:** The help information for all of the options within the **UEFI RAID Configuration Utility** can be viewed by clicking on **Help** at the top right-hand corner of the browser screen. Help information for individual option menus can also be viewed by scrolling down on each option.

 **NOTE:** Some of the options within the **UEFI RAID Configuration Utility** are not present if the controller does not support the corresponding feature. Options may also be grayed out if the feature is supported in existing configuration.

 **NOTE:** Background operations are blocked on the PERC H310 controller in the **UEFI RAID Configuration Utility** and no operation progress updates are to be seen.

Exiting The UEFI RAID Configuration Utility

To exit the **UEFI RAID Configuration Utility**, press **<Esc>** or click on **Back** at any browser screen, until you return to the **Device Settings** menu.

Configuration Options

There are four main configuration option menus available in the **UEFI RAID Configuration Utility**. These menus are outlined below.

- **Controller Management**—Performs controller properties, foreign configurations, and battery, and other high-level controller functions.
- **Virtual Disk Management**—Creates or deletes virtual disks, manage virtual disk properties, and performs operations such as background initialization, consistency check, locate and reconfiguration.
- **Physical Disk Management**—Displays physical disk properties and performs operations such as assign or unassign hot spares, locate, force online/offline and rebuild after a physical disk failure.
- **Enclosure Management**—Displays attached enclosures, enclosure firmware revision and physical disks attached to the enclosures.

Controller Management Menu

The **Controller Management** menu can be used to perform the following controller-level functions. Each of these functions is its own selectable menu option.

- **View Controller Information**—Displays information about the controller.
- **Change Controller Properties**—Updates controller properties and/or restores factory defaults for the controller.
 - PCI link speed: For systems capable of PCI Generation 3 link speeds, an option will appear at the top of this screen named **Change Link Speed to Gen 3**. When selected, the system will need to be rebooted. If the system is currently using Gen 3, the option will read **Change Link Speed to Gen 2**.
 - Boot Error Handling: Selects the way errors found at boot are handled.
 - * Stop on errors: Interrupt the boot on all errors. Requires the user to either continue boot or enter the BIOS Configuration Utility to resolve errors.
 - * Pause on errors: Interrupt the boot on serious errors. Requires the user to either continue boot or enter the BIOS Configuration Utility to resolve errors.
 - * Ignore errors: The controller will take the default action for each error and will only halt boot when continuing is impossible.

- * **Safe mode on errors:** The system will continue booting as long as the boot volume is not on this controller. Drives connected to the controller will be reported as unconfigured good drives to the RAID management utilities upon a successful boot. No configuration changes will be allowed, except to clear the controller configuration if desired. Rebooting the system will exit safe mode, and drives will appear as foreign on the next boot; if problems remain the system will enter safe mode again. If the boot volume is on this controller, boot will halt.

- **Battery Management**—Displays battery status and capacity information.
- **Clear Configuration** — Deletes all existing configurations on the controller.
- **Manage Foreign Configuration**—Displays import and/or clears foreign configurations.
- **Save Controller Events**—Saves a copy of the controller events file in the desired directory.
- **Clear Controller Events**—Deletes all existing entries in the controller events file.
- **Save Debug Log**—Saves a copy of the firmware’s terminal log entries for the controller.
- **Enable Security** — Enables security on the controller via the desired security mode.
- **Disable Security**—Disables security on the controller.
- **Change Security Key**—Changes the security key or switches between security modes on the controller.
- **Convert To RAID Capable**—Allows conversion of physical disks from Non-RAID to RAID Capable.
- **Convert To Non-RAID**—Allows conversion of physical disks from RAID Capable to Non-RAID.

Virtual Disk Management

The **Virtual Disk Management** menu can be used to create and manage virtual disks on the controller. The following functions can be performed under **Virtual Disk Management**. Each of these functions is its own selectable menu option.

- **Create Virtual Disk**—Creates a virtual disk by selecting the RAID level, physical disks, and virtual disk parameters.
- **Manage Virtual Disk Properties**—Displays and manages virtual disk properties, view associated physical disks and manage dedicated hot spares.
- **Select Virtual Disk Operations**—Selects and executes virtual disk operations such as delete, locate, secure, expand, initialization and consistency check.
- **View Disk Group Properties**—Displays information about the available disk groups, associated virtual disks, and the capacity allocation.
- **Reconfigure Virtual Disks**—Performs a RAID Level Migration (RLM) which converts a virtual disk to a different RAID level.
- **Create CacheCade Virtual Disks**—Creates a CacheCade virtual disk by selecting solid state disks (SSDs).

Physical Disk Management Menu

The **Physical Disk Management** menu can be used to view and manage physical disks on the controller. The following functions can be performed under **Physical Disk Management**. Each of these functions is its own selectable menu option.


- **View Physical Disk Properties**—Displays information about all the physical disks including their state.
- **Select Physical Disk Operations**—Selects and execute physical disk operations such as locate, assign/unassign hot spare, replace member, force offline/online, and convert to RAID Capable/Non-RAID.
- **View Global Hot Spares**—Displays all the assigned global hot spares on the controller.

Enclosure Management

The **Enclosure Management** menu can be used to view the firmware version of the enclosure and the physical disks attached to the controller, along with their physical disk states.

CacheCade

The Dell PowerEdge RAID Controller (PERC) H710, H710P, and H810 cards support CacheCade, a feature that can improve application performance by increasing read caching capacity. The CacheCade feature makes use of high performing solid state disks (SSDs) as a secondary tier of cache. CacheCade provides faster reads and maximizes transactional I/O performance.


 **NOTE:** PERC H310 does not support Cachecade.

The use of SSDs for caching allows a large quantity of data to be present in the cache, resulting in performance improvement in read-intensive applications. Some examples of read-intensive applications include online transaction processing (OLTP), file server, and web server workloads. CacheCade allows for an increase in the I/O performance of hard disk drive (HDD)-based disk groups with the assistance of SSD technology.

CacheCade Virtual Disk Characteristics

The following are the CacheCade virtual disk characteristics:

- Support for CacheCade virtual disks exists only on controllers containing 1 GB of Non-Volatile (NV) Cache.
- CacheCade virtual disks can only be created with SSDs.
- The maximum combined size of CacheCade virtual disks is 512 GB.

 **NOTE:** Multiple CacheCade virtual disks may be created, but they are combined to operate as a single cache pool up to the maximum size.

- Data on virtual disks containing secured Self-Encrypting Disks (SEDs) or SSDs will not be cached by CacheCade.

 **NOTE:** Data on VD's with unsecured SEDs can be cached using CacheCade.

- CacheCade virtual disks only cache input reads that are smaller than 64 KB.
- CacheCade virtual disks are read cache only.
- CacheCade virtual disks cannot be migrated to a controller that does not support CacheCade.
- Importing a CacheCade drive may result in a RAID 0 VD. You must reconfigure the CacheCade VD after importing.
- CacheCade virtual disks are not presented to the operating system.

Configuring And Managing CacheCade Virtual Disks

All management applications allow the creation and deletion of CacheCade virtual disks.

The following sections describe the menu options specific to CacheCade virtual disk management and provide detailed instructions to perform the configuration tasks. The contents of the following section apply to the **BIOS Configuration Utility** (<Ctrl> <R>). For more information on the management applications, see the topic [Management Applications For PERC Cards](#).

CacheCade Virtual Disk Management

The Virtual Disk Management screen is the first screen that is displayed when you access a RAID controller from the main menu screen on the BIOS Configuration Utility.

The following are CacheCade-related actions you can perform through the virtual disk management menu:

- Create CacheCade virtual disk
- Assign CacheCade virtual disk name
- Delete virtual disk

Creating CacheCade Virtual Disks



NOTE: Only SSDs can be used to create CacheCade virtual disks.





NOTE: Combining SAS and SATA SSDs within a CacheCade virtual disk is not supported.



NOTE: To avoid inefficient use of space, it is recommended that CacheCade volumes are created with drives of the same size. Larger disk sizes are truncated to the size of the smallest contributing disk, similar to RAID 0.


To create a CacheCade virtual disk:


1. During host system bootup, press **<Ctrl> <R>** when the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed. If there are more than one controller, the main menu screen is displayed.
2. Select a controller, and press **<Enter>**.
The **Virtual Disk Management** screen is displayed for the selected controller.
3. Use the arrow keys to highlight **Controller #**.
4. Press **<F2>**.
The list of available actions is displayed.
5. Select **Create CacheCade Virtual Disk** and press **<Enter>**.
The **Create CacheCade Virtual Disk** screen is displayed. The cursor is on the first SSD listed in the **Select SSD** section.
6. Select the desired SSD(s). As each new SSD is selected, the **CacheCade Virtual Disk Size** changes to reflect the new size.
 **NOTE:** You cannot specify the size of the virtual disk.
7. Press **<Tab>** to move the cursor to the **CacheCade Virtual Disk Name** field. Enter a name if required.
8. After you specify the virtual disk name, select **OK** to save the selection or select **Cancel** to cancel the selection.
After the CacheCade virtual disk is created successfully, it is listed in the **Virtual Disk Management** screen under the CacheCade disk group, and is labeled as a CacheCade virtual disk. The virtual disk has an optimal state and its RAID level is RAID 0.
 **NOTE:** You can only delete or rename a CacheCade virtual disk. Background initialization, fast initialization, full initialization, and consistency check operations are not applicable to CacheCade virtual disks.

Deleting CacheCade Virtual Disks

To delete CacheCade virtual disks, perform the following steps in the **BIOS Configuration Utility**:

1. Press <Ctrl> <N> to access the **Virtual Disk Management** screen.
2. Use the arrow keys to move the cursor to the **CacheCade Disk Group** or **Virtual Disks** heading.
3. Press <F2>.
The **Action** menu is displayed.
4. Select **Delete VD** and press <Enter>.

 **NOTE:** Warning messages are displayed stating the effect of deleting a virtual disk. Click **OK** to complete the virtual disk deletion.

 **NOTE:** In operating system management applications, a CacheCade virtual disk deletion or removal can be done without interrupting any I/Os. The controller stops caching via the secondary cache, but all outstanding I/Os are completed.

Reconfiguring CacheCade Virtual Disks

In operating system management applications, the resizing of CacheCade virtual disks occurs without stopping any existing I/Os. Creating a new CacheCade virtual disk or adding one or more SSDs to an existing CacheCade virtual disk increases the total cache size. The new resources are immediately used after the addition.

There are two methods to reconfigure CacheCade virtual disks:


- Automatic reconfiguration of cachecade virtual disks
- Manual reconfiguration of cachecade virtual disks


Automatic Reconfiguration Of CacheCade Virtual Disks

A CacheCade virtual disk that is made up of more than one SSD is automatically reconfigured upon a removal or failure of a member SSD. The virtual disk retains an Optimal state and adjusts its size to reflect the remaining number of member disks. If auto-rebuild is enabled on the controller, when a previously removed SSD is inserted back into the system or replaced with a new compatible SSD, the CacheCade automatically reconfigures and adjusts its size to reflect the addition of the member SSD.

The number of SSDs to be removed from a CacheCade virtual disk cannot equal the total number of SSDs currently in the CacheCade virtual disk.

After the automatic reconfiguration and resizing of a CacheCade virtual disk, the new virtual disk size is displayed in both the **BIOS Configuration Utility** as well as in the OpenManage storage management application.

 **CAUTION:** If a disk is removed from a CacheCade virtual disk, the associated slot for that disk is a hot slot for the CacheCade volume. Any compatible disk inserted into that slot is automatically added to the CacheCade virtual disk. Any preexisting data on that drive is lost.

 **NOTE:** Disks inserted into a CacheCade hot slot must be equal to or greater than the smallest contributing disk in the virtual disk.

Manual Resizing Of CacheCade Virtual Disks


In operating system management applications, a manual resizing of a CacheCade virtual disk occurs due to the addition or removal of drives. Reboot is not necessary.

- Any number of SSDs can be added to a CacheCade virtual disk.
- There is no SAS and SATA mixing allowed within a CacheCade virtual disk, so SATA SSDs cannot be added to a SAS CacheCade virtual disk and vice versa.
- HDDs cannot be added to a CacheCade virtual disk.



NOTE: Capacity of drives added to a CacheCade volume must be equal to or greater than the smallest contributing drive in the virtual disk. The manual resizing of a CacheCade virtual disk cannot be initiated in the **BIOS Configuration Utility**. It can only be initiated in the OpenManage storage management application.

Security Key And RAID Management

 **NOTE:** PERC H310 does not support any security implementations.

Security Key Implementation

The Dell PowerEdge RAID Controller (PERC) H710, H710P, and H810 cards support Self-Encrypting Disks (SED) for protection of data against loss or theft of SEDs. Protection is achieved by the use of encryption technology on the drives. There is one security key per controller. You can manage the security key under Local Key Management (LKM). The key can be escrowed in to a file using Dell OpenManage. The security key is used by the controller to lock and unlock access to encryption-capable physical disks. In order to take advantage of this feature, you must:

1. Have SEDs in your system.
2. Create (LKM) a security key.

Security Key Management In The BIOS Configuration Utility


The Dell OpenManage storage management application and the **BIOS Configuration Utility** (<Ctrl> <R>) of the controller allow security keys to be created and managed as well as create secured virtual disks. The following section describes the menu options specific to security key management and provide detailed instructions to perform the configuration tasks. The contents of the following section apply to the **BIOS Configuration Utility** (<Ctrl> <R>). For more information on the management applications, see the topic [Management Applications For PERC Cards](#).

- The Virtual Disk Management screen, **VD Mgmt**, is the first screen that is displayed when you access a RAID controller from the main menu screen on the **BIOS Configuration Utility** (<Ctrl> <R>). The following are security-related actions you can perform through the virtual disk management menu:
 - **Security Key Management**—Creates, changes, or deletes the security settings on a controller.
 - **Secure Disk Group**—Secures all Virtual Disks in Disk Group.
- The Physical Disk Management screen, **PD Mgmt**, displays physical disk information and action menus. The following are security-related actions you can perform through the physical disk management menu:
 - **Instant Secure Erase**—Permanently erases all data on an encryption-capable physical disk and resets the security attributes.

For more information on the Physical Disk Management screen and Virtual Disk Management screen, see the topics [Physical Disk Management \(PD Mgmt\)](#) and [Virtual Disk Management](#) respectively.

Local Key Management (LKM)

You can use LKM to generate the key ID and the passphrase required to secure the virtual disk. You can secure virtual disks, change security keys and manage secured foreign configurations using this security mode.

 **NOTE:** Under LKM, you are prompted for a passphrase when you create the key.

Creating A Security Key

Perform the following steps to create a security key on the controller:



NOTE: There is no passphrase backup option when you create a security key; you need to remember your passphrase.

1. During the host system boot up, press **<Ctrl> <R>** when the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed.
If there is more than one controller, the main menu screen is displayed.
2. Select a controller, and press **<Enter>**.
The **Virtual Disk Management** screen is displayed for the selected controller.
3. Use the arrow keys to highlight **Security Key Management**.
4. Press **<F2>** to display the actions you can perform.
5. Select **Create Key** and press **<Enter>**.
The **Create Security Key** screen is displayed. The cursor is at the **Security Key Identifier**.
6. Enter an identifier for your security key.



NOTE: The Security Key Identifier is a user supplied clear text label used to associate the correct security key with the controller.

7. Press **<Tab>** to enter a passphrase.



CAUTION: If you lose or forget your passphrase, the data in your virtual disk is inaccessible.

8. Press **<Tab>** and select **OK** to accept the settings and to exit the window. Select **Cancel** to exit if you do not want to create a security key on the controller.

Changing The Security Key



NOTE: Change Key is active if there is a security key present on the controller.


Perform the following steps when changing the security key on the controller:

1. During host system bootup, press **<Ctrl> <R>** when the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed.
If there is more than one controller, the main menu screen is displayed.
2. Select a controller, and press **<Enter>**.
The **Virtual Disk Management** screen is displayed for the selected controller.
3. Use the arrow keys to highlight **Security Key Management**.
4. Press **<F2>** to display the actions you can perform.
5. Select **Change Key** and press **<Enter>**.
6. The **Change Security Key** screen is displayed. The cursor is at the **Security Key Identifier**. Enter an identifier for your Security Key.
7. Press **<Tab>** to enter a new passphrase.





NOTE: You need to provide the current passphrase in order to change the security key on the controller.


8. Press **<Tab>** and select **OK** to accept the settings and to exit the window. Select **Cancel** to exit if you do not want to change the security key on the controller.

 **NOTE:** If there is an existing configuration on the controller, it is updated with the new security key. If you had previously removed any secured disks, you still need to supply the old passphrase to import them.

Deleting A Security Key

 **NOTE:** Delete Key is active if there is a security key present on the controller.

 **NOTE:** Delete Key can only be performed when there are no secured virtual disks present.


 **NOTE:** After the Delete Key operation, all unconfigured, secured SEDs are secure-erased.

Perform the following steps when deleting the security key on the controller:


1. During host system bootup, press **<Ctrl> <R>** when the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed. If there is more than one controller, the main menu screen is displayed.
2. Select a controller, and press **<Enter>**.
The **Virtual Disk Management** screen is displayed for the selected controller.
3. Use the arrow keys to highlight **Security Key Management**.
4. Press **<F2>** to display the actions you can perform.
5. Select the **Delete Key** and press **<Enter>**.

Creating Secured Virtual Disks

To create a secured virtual disk, the controller must have a security key established first. See the topic [Creating A Security Key](#).

 **NOTE:** Combining SAS and SATA hard drives within a virtual disk is not supported. Also, combining hard drives and solid state drives (SSDs) within a virtual disk is not supported.

After the security key is established, perform the steps outlined in the topic [Creating Virtual Disks](#) to create a virtual disk. To secure the virtual disk, navigate to the **Secure VD** option at the bottom left area of the **Create New VD** screen.

 **NOTE:** All virtual disks added to a secured Disk Group are secured.

Securing Pre-Existing Virtual Disks

If an unsecured virtual disk was created on a controller, you can secure the virtual disk as long as the following conditions are met:

- The controller has a security key present.
- All the physical disks in the virtual disk are SED.

Perform the following steps when securing a pre-existing virtual disk on a secured controller:

1. During the host system bootup, press **<Ctrl> <R>** when the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed.
If there is more than one controller, the main menu screen is displayed.
2. Select a controller, and press **<Enter>**.
The **Virtual Disk Management** screen is displayed for the selected controller.
3. Use the arrow keys to highlight the **Disk Group** number.
4. Press **<F2>** to display a menu of the available actions.

5. Highlight the **Secure Disk Group** option and press **<Enter>**.

 **NOTE:** If you select to secure a Disk Group, all VD's part of the Disk Group are secured.

Securing Pre-Existing Virtual Disks

If an unsecured virtual disk was created on a controller, you can secure the virtual disk as long as the following conditions are met:

- The controller has a security key present.
- All the physical disks in the virtual disk are SED.


Perform the following steps when securing a pre-existing virtual disk on a secured controller:


1. During the host system bootup, press **<Ctrl> <R>** when the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed.
If there is more than one controller, the main menu screen is displayed.
2. Select a controller, and press **<Enter>**.
The **Virtual Disk Management** screen is displayed for the selected controller.
3. Use the arrow keys to highlight the **Disk Group** number.
4. Press **<F2>** to display a menu of the available actions.
5. Highlight the **Secure Disk Group** option and press **<Enter>**.


 **NOTE:** If you select to secure a Disk Group, all VD's part of the Disk Group are secured.


Importing Or Clearing Secured Foreign Configurations And Secure Disk Migration


Secured virtual disks created on a PERC H700/H800 or H710/H710P/H810 card can be migrated to another PERC H710, H710P, or H810 card. A virtual disk secured with a security key different from the current controller security key cannot be imported without authentication of the original passphrase used to secure them. When importing secured virtual disk(s) created with a different security key, the secured foreign configuration(s) do not show in the **Foreign Configuration View** screen. Follow the steps below to import or clear a foreign secured virtual disk.

 **NOTE:** If you are importing secured and unsecured virtual disks, you are prompted to resolve the secured foreign configuration first.

 **NOTE:** The PERC H710, H710P, or H810 card needs to have a security key present before being able to import a secured virtual disk.

 **NOTE:** Any unsecured virtual disks imported are still unsecured.


 **NOTE:** If you are importing a virtual disk originally secured with a local key (LKM), you are prompted for the passphrase used to secure that virtual disk.


 **NOTE:** A secured VD cannot be imported using the PERC H310 card.

Perform the following steps when importing a foreign secured virtual disk:

1. During the host system bootup, press **<Ctrl> <R>** when the BIOS screen is displayed.
The **Virtual Disk Management** screen is displayed.
If there is more than one controller, the main menu screen is displayed.
2. Select a controller, and press **<Enter>**.
The **Virtual Disk Management** screen is displayed for the selected controller.

3. Press <F2> to display a menu of available actions.
4. Select **Import** to import the foreign configuration or **Clear** to delete the foreign configuration. Press <Enter>.

 **NOTE:** To **Clear**, you need to **Instant Secure Erase** foreign configurations secured with a different security key.
5. If you select to **Import** the configuration, the **Secure Foreign Import** screen is displayed.

 **NOTE:** The key identifier for the passphrase used to secure the foreign secured virtual disks is displayed under the **Secured Drives** option.
6. Enter the passphrase that was used to secure the foreign configuration.
7. Press <Tab> and select **OK** to finish importing the secured foreign configuration or select **Cancel** to exit this menu.

If you select **Cancel** for the secured foreign import, the disks remain inaccessible until imported or instant secure erased. See the topic [Instant Secure Erase](#).

Instant Secure Erase

Instant Secure Erase is the process of permanently erasing all data on an encryption-capable physical disk and resetting the security attributes. You need to execute **Instant Secure Erase** on SEDs that are inaccessible (blocked) due to a lost or forgotten passphrase.

 **CAUTION:** By executing **Instant Secure Erase**, the data on your ☐ encryption-capable physical disk is lost.

To execute Instant Secure Erase:

1. Press <Ctrl> <N> to access the **PD Mgmt** screen.

A list of physical disks is displayed. On the right menu, the physical disk properties are displayed including information about whether the physical disk is secured or not.
2. Press the down-arrow key to highlight a physical disk that is secured.
3. Press <F2> to display a menu of available actions.
4. The **Secure Erase** option is highlighted at the bottom of the menu.
5. Press <Enter> to **Secure Erase** the physical disk and select **YES**.

Troubleshooting Security Key Errors

Secured Foreign Import Errors

A foreign configuration is a RAID configuration that already exists on a replacement physical disk that you install in a system. A secured foreign configuration is a RAID configuration that was created under a different security key.

There are two scenarios in which a secured foreign import fails:

- **The passphrase authentication fails**—A VD secured with a security key different from the current controller security key cannot be imported without authentication of the original passphrase used to secure them. Supply the correct passphrase to import the secured foreign configuration. If you have lost or forgotten the passphrase, the secured foreign disks remain locked (inaccessible) until the appropriate passphrase is entered or if they are instant secure erased.
- **The secured VD is in an offline state after supplying the correct passphrase**—You must check to determine why the virtual disk failed and correct the problem. See the topic [Troubleshooting](#).

Failure to Select Or Configure Non Self-Encrypting Disks (Non-SED)

A virtual disk can be either secured or unsecured depending on how it was configured when created. In order to create a secured virtual disk, the controller must have a security key present and must be composed of SEDs only. In order to select/configure non-SED, you must create an unsecured virtual disk. You can create an unsecured virtual disk even if

there is a security key present. Select the **Secure VD** option as **No** in the **Create New VD** menu. See the topic [Creating Virtual Disks](#) for steps on how to create an unsecured virtual disk.

Failure To Delete Security Key

A security key is used to lock or unlock access to a security-enabled component. This key is not utilized in the actual encryption of data. If a security key is present, both secured and unsecured virtual disks may exist.

To delete the security key, you must have a previously established security key present on the controller and there cannot be any configured secured disks. If there are configured secured disks, remove or delete them.

Failure To Instant Secure Erase Task On Physical Disks

Instant Secure Erase is the process of securely erasing all data permanently on an encryption-capable physical disk and resetting the security attributes. It is used in a couple of scenarios such as deleting a foreign configuration in the event of a forgotten or lost passphrase or unlocking a disk that had been previously locked.

Instant Secure Erase can be executed only on encryption-capable disks as long as the disks are not hot spares and are not configured (part of a virtual disk). Ensure that the conditions are met and see the topic [Instant Secure Erase](#).

Troubleshooting

To get help with your Dell PowerEdge RAID Controller (PERC) H310, H710, H710P and H810, you can contact your Dell Technical Service representative or see dell.com/support.

BIOS Error Messages

The controller BIOS read-only memory (ROM) provides Int 13h functionality (disk I/O) for the virtual disks connected to the controller. You can boot from or access the physical disks without a driver.

Discovery Error Message

Error Message:	A discovery error has occurred, please power cycle the system and all the enclosures attached to this system.
Probable Cause:	This message indicates that discovery did not complete within 120 seconds. The SAS cables for your system might be improperly connected.
Corrective Action:	Check the cable connections and fix any problems. Restart the system.

Extra Enclosure Error Message

Error Message:	There are X enclosures connected to connector Y, but only maximum of 4 enclosures can be connected to a single SAS connector. Please remove the extra enclosures then restart your system.
Probable Cause:	This message is displayed when the BIOS detects more than four enclosures connected to a single SAS connector.
Corrective Action:	You must remove all additional enclosures and restart your system.

Cache Data Lost Error Message

Error Message:	Cache data was lost, but the controller has recovered. This could be due to the fact that your controller had protected cache after an unexpected power loss and your system was without power longer than the battery backup time. Press any key to continue or 'C' to load the configuration utility.
Probable Cause:	This message is displayed under the following conditions: <ul style="list-style-type: none"> The adapter detects that the cache in the controller cache has not yet been written to the disk subsystem.

- The controller detects an Error-Correcting Code (ECC) error while performing its cache checking routine during initialization.
- The controller discards the cache rather than sending it to the disk subsystem because the data integrity cannot be guaranteed.

Corrective Action: To resolve this issue, allow the battery to charge fully. If the problem persists, the battery or controller memory may be faulty; contact Dell Technical Support.

Missing Disks In Virtual Disk Error Message

Error Message: The following virtual disks have missing disks: (x). If you proceed (or load the configuration utility), these virtual disks will be marked OFFLINE and will be inaccessible. Please check your cables and ensure all disks are present. Press any key to continue, or 'C' to load the configuration utility.

Probable Cause: The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.

Corrective Action: Check the cable connections and fix any issues. Restart the system. If there are no cable problems, press any key or <C> to continue.

Previous Configuration Of Disks Removed Error Message

Error Message: All of the disks from your previous configuration are gone. If this is an unexpected message, then please power off your system and check your cables to ensure all disks are present. Press any key to continue, or press <C> to load the configuration utility.

Probable Cause: The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.

Corrective Action: Check the cable connections and fix any issues before restarting the system. If there are no cable problems, press any key or <C> to continue.

Missing Virtual Disks Error Message

Error Message: The following virtual disks are missing: (x). If you proceed (or load the configuration utility), these virtual disks will be removed from your configuration. If you wish to use them at a later time, they will have to be imported. If you believe these virtual disks should be present, please power off your system and check your cables to ensure all disks are present. Press any key to continue, or 'C' to load the configuration utility.

Probable Cause: The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.

Corrective Action:	Check the cable connections and fix any issues. Restart the system. If there are no cable problems, press any key or <C> to continue.
---------------------------	---

Dirty Cache Data Error Message

Error Message:	The following virtual disks are missing: (x). If you proceed (or load the configuration utility), these virtual disks will be removed from your configuration. If you wish to use them at a later time, they will have to be imported. If you believe these virtual disks should be present, please power off your system and check your cables to ensure all disks are present. Press any key to continue, or 'C' to load the configuration utility. The cache contains dirty data, but some virtual disks are missing or will go offline, so the cached data cannot be written to disk. If this is an unexpected error, then please power off your system and check your cables to ensure all disks are present. If you continue, the data in cache will be permanently discarded. Press 'X' to acknowledge and permanently destroy the cached data.
Probable Cause:	The controller preserves the dirty cache from a virtual disk if the disk becomes offline or is deleted because of missing physical disks. This message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.
Corrective Action:	Check the cable connections and fix any problems. Restart the system. Use the <Ctrl> <R> utility to import the virtual disk or discard the preserved cache. For the steps to manage preserved cache, see the topic Managing Preserved Cache .

BIOS Disabled Error Message

Error Message:	BIOS Disabled. No Logical Drives Handled by BIOS.
Probable Cause:	This warning message is displayed after you disable the ROM option in the configuration utility. When the ROM option is disabled, the BIOS cannot boot to Int 13h and cannot provide the ability to boot from the virtual disk. Int 13h is an interrupt signal that supports numerous commands that are sent to the BIOS and then passed to the physical disk. The commands include actions you can perform with a physical disk, such as reading, writing, and formatting.
Corrective Action:	Enable the ROM option.

Drive Configuration Changes Error Message

Error Message:	Entering the configuration utility in this state will result in drive configuration changes. Press 'Y' to continue loading the configuration utility or please power off your system and check your cables to ensure all disks are present and reboot.
Probable Cause:	The message is displayed after another BIOS warning indicating there are problems with previously configured disks and you have chosen to accept any changes and continue. The SAS cables for your system might be improperly connected.

Corrective Action: Check the cable connections and fix any problems before restarting the system. If there are no cable problems, press any key or <Y> to continue.

Adapter At Baseport Not Responding Error Message

Error Message: Adapter at Baseport xxxx is not responding, where xxxx is the baseport of the controller.

Corrective Action: Contact Dell Technical Support.

Offline Or Missing Virtual Drives With Preserved Cache Error Message

Error Message: There are offline or missing virtual drives with preserved cache. Please check the cables and ensure that all drives are present. Press any key to enter the configuration utility.

Probable Cause: The controller preserves the dirty cache from a virtual disk if the disk becomes offline or is deleted because of missing physical disks. This preserved dirty cache is called pinned cache, and is preserved until you import the virtual disk, or discard the cache.

Corrective Action: Use the <Ctrl> <R> utility to import the virtual disk or discard the preserved cache. For the steps used to manage preserved cache, see the topic [Managing Preserved Cache](#).

Virtual Disks Offline Error Message

Error Message: x Virtual Disk(s) Degraded, where x is the number of virtual disks degraded.

Probable Cause: This message is displayed when the BIOS detects virtual disks in a degraded state.

Corrective Action: Take corrective action(s) to make the virtual disks optimal. The BIOS does not take any action.

Virtual Disks Degraded Error Message

Error Message: x Virtual Disk(s) Degraded, where x is the number of virtual disks degraded.

Probable Cause: This message is displayed when the BIOS detects virtual disks in a degraded state.

Corrective Action: To make the virtual disks optimal, take one the following corrective actions:

- Ensure all disks in the Virtual Disk are present and online.
- Replace any failed disks that may be in the array.
- Correct a hot spare disk, and rebuild the array.

The BIOS does not take any action.

Virtual Disks Partially Degraded Error Message

Error Message:	x Virtual Disk(s) Partially Degraded, where x is the number of virtual disks partially degraded.
Probable Cause:	This message is displayed when the BIOS detects a single disk failure in a RAID 6 or RAID 60 configuration.
Corrective Action:	You must check why the member disk is not present to correct the problem. The BIOS does not take any action.

Memory Or Battery Problem Error Message

Error Message:	Memory/Battery problems were detected. The adapter has recovered, but cached data was lost. Press any key to continue
Probable Cause:	The message occurs under the following conditions: <ul style="list-style-type: none">• The adapter detects data in the controller cache that has not yet been written to the disk subsystem.• The controller detects an Error-Correcting Code (ECC) error while performing its cache checking routine during initialization.• The controller discards the cache rather than sending it to the disk subsystem because the data integrity cannot be guaranteed.• The battery may be under charged.
Corrective Action:	Allow the battery to charge fully to resolve this problem. If the problem persists, the battery or controller memory may be faulty. Contact Dell Technical Support.

Firmware Fault State Error Message

Error Message:	Firmware is in Fault State.
Corrective Action:	Contact Dell Technical Support.

Foreign Configuration Found Error Message

Error Message:	Foreign configuration(s) found on adapter. Press any key to continue, or 'C' to load the configuration utility or 'F' to import foreign configuration(s) and continue.
Probable Cause:	When a controller firmware detects a physical disk with existing foreign metadata, it flags the physical disk as <i>foreign</i> and generates an alert indicating that a foreign disk was detected.
Corrective Action:	Press <F> at this prompt to import the configuration (if all member disks of the virtual disk are present) without loading the BIOS Configuration Utility (<Ctrl> <R>). Or press <C> to enter the BIOS Configuration Utility (<Ctrl> <R>) and either import or clear the foreign configuration.

Foreign Configuration Not Found In <Ctrl> <R> Error Message

Error Message: The foreign configuration message is present during POST but no foreign configurations are present in the foreign view page in <Ctrl> <R>. All virtual disks are in an optimal state.

Corrective Action: **Corrective Action:** Ensure all your PDs are present and all VDs are in optimal state. Clear the foreign configuration using <Ctrl> <R> or **Dell OpenManage Server Administrator Storage Management**.



CAUTION: The physical disk goes to Ready state when you clear the foreign configuration.

If you insert a physical disk that was previously a member of a virtual disk in the system, and that disk's previous location has been taken by a replacement disk through a rebuild, you must manually remove the foreign configuration flag of the newly inserted disk.

Previous Configuration Cleared Or Missing Error Message

Error Message: Previous configuration(s) cleared or missing. Importing configuration created on XX/XX XX.XX. Press any key to continue, or 'C' to load the configuration utility.

Probable Cause: The message indicates that the controller and physical disks have different configurations.

Corrective Action: You can use the **BIOS Configuration Utility** (<Ctrl> <R>) to clear the foreign configuration.


Invalid SAS Topology Detected Error Message

Error Message: Invalid SAS topology detected. Please check your cable configurations, repair the problem, and restart your system.

Probable Cause: The SAS cables for your system are improperly connected.

Corrective Action: Check the cable connections and fix problems if any. Restart the system.

Multibit ECC Errors Detected Error Messages

Error Message:  **NOTE:** The following error is specific to the PERC H710 controller only.

Multibit ECC errors were detected on the RAID controller. If you continue, data corruption can occur. Contact technical support to resolve this issue. Press 'X' to continue or else power off the system, replace the controller and reboot.

Or

Multibit ECC errors were detected on the RAID controller. The DIMM on the controller needs replacement. If you continue, data corruption can occur. Press 'X' to continue or else power off the system, replace the DIMM module, and reboot. If you have replaced the DIMM please press 'X' to continue.



CAUTION: When removing or replacing the PERC Mini Monolithic Controller, hold the card by its edges marked by the blue touch points. Do not handle the card while holding the battery or the heatsink.

For more information, see figure “Removing and Installing the PERC Mini Monolithic Controller” in [Removing The PERC Controller](#).



CAUTION: Multi-bit ECC errors cause data corruption and data loss. In case of MBE errors, contact Dell Technical Support.

Probable Cause: Multi-bit ECC errors occur in the memory and can corrupt cached data and discard it.

Corrective Action: Contact Dell Technical Support.

Configured Disks Removed Or Not Accessible Error Message

Error Message: Some configured disks have been removed from your system or are no longer accessible. Check your cables and ensure all disks are present. Press any key or 'C' to continue.

Probable Cause: The message indicates that some configured disks were removed. If the disks were not removed, they are no longer accessible. The SAS cables for your system might be improperly connected.

Corrective Action: Check the cable connections and fix any issues. Restart the system. If there are no cable problems, press any key or <C> to continue.

Battery Discharged Or Disconnected Error Message

Error Message: The battery is currently discharged or disconnected. Verify the connection and allow 30 minutes for charging. If the battery is properly connected and it has not returned to operational state after 30 minutes of charging then contact technical support for additional assistance.

Probable Cause: This message may appear for controller without battery, depending on the virtual disks' policies.

Corrective Action:

- The controller battery is missing or currently inaccessible, contact Dell support if the problem persist after 30 minutes.
- The controller battery is completely discharged and needs to be charged for it to become active. You must first charge the battery and wait for a few minutes for the system to detect it.

Degraded State Of Virtual Disks

A redundant virtual disk is in a degraded state when one or more physical disks have failed or are inaccessible. For example, if a RAID 1 virtual disk consists of two physical disks and one of them fails or become inaccessible, the virtual disk become degraded.


To recover a virtual disk from a degraded state, you must replace the failed physical disk and rebuild it. Once the rebuilding process is complete, the virtual disk state changes from degraded to optimal. For information on rebuilding the disk, see the topic [Performing A Manual Rebuild Of An Individual Physical Disk](#).

Memory Errors

Memory errors can corrupt cached data, so the controllers are designed to detect and attempt to recover from the memory errors. Single-bit memory errors can be handled by the controller and do not disrupt normal operation. A notification is sent if the number of single-bit errors exceeds a threshold value.

Multi-bit errors are more serious as they result in corrupted data and data loss. The following are the actions that occur in the case of multi-bit errors:

- If a multi-bit error occurs while accessing data in the cache when the controller is started with dirty cache, the controller discards the cache contents. The controller generates a warning message to the system console to indicate that the cache was discarded and generates an event.
- If a multi-bit error occurs at run-time either in code/data or in the cache, the controller stops.
- The controller logs an event to the controller's internal event log and a message during POST is displayed indicating a multi-bit error has occurred.

 **NOTE:** In case of a multi-bit error, contact Dell Technical Support.

Preserved Cache State

The controller preserves the dirty cache from a virtual disk if the virtual disk becomes offline or is deleted because of missing physical disks. This preserved dirty cache is called pinned cache and is preserved until you import the virtual disk or discard the cache.

Use the **BIOS Configuration Utility** (<Ctrl> <R>) to select whether to import the virtual disk or discard the preserved cache. In the **VD Mgmt** menu, select **Manage Preserved Cache** and follow the steps on the screen.

General Issues

PERC Card Has Yellow Bang In Device Manager

Issue:	The device is displayed in Device Manager but has a yellow bang (exclamation mark).
Corrective Action:	Reinstall the driver. For more information on reinstalling drivers, see the topic Driver Installation .

PERC Card Not Seen In Device Manager

Issue:	The device does not appear in the Device Manager .
Correction:	Turn off the system and reseal the controller. For more information see, Installing The PERC Controller .

No Hard Drives Found Error Message During Microsoft Windows Server 2003 Installation

Issue:	No Hard Drives Found message is displayed during a media-based installation of Microsoft Windows Server 2003.
---------------	---

**Corrective
Action:**

**The issue occurs
because of any of
the following
causes:**

- The driver is not native in the operating system.
To resolve the issue, press <F6> to install the RAID device driver during installation.
- The virtual disks are not configured properly.
To resolve the issue, enter the **BIOS Configuration Utility** (<Ctrl> <R>) to configure the virtual disks. For procedures to configure the virtual disks, see the topic [Management Applications For PERC Cards](#).
- The controller BIOS is disabled.
To resolve the issue, enter the **BIOS Configuration Utility** (<Ctrl> <R>) to enable the BIOS. For information on configuring virtual disks, see the topic [Management Applications For PERC Cards](#).

Physical Disk Issues

Physical Disk In Failed State

Issue: One of the physical disks in the disk array is in the failed state.

**Corrective
Action:** Update the PERC cards to the latest firmware available on dell.com/support.

Unable to Rebuild A Fault Tolerant Virtual Disk

Issue: Cannot rebuild a fault tolerant virtual disk. For more information, see the alert log for virtual disks.

**Corrective
Action:** The replacement disk is too small or not compatible with the virtual disk. Replace the failed disk with a compatible good physical disk with equal or greater capacity.

Fatal Error Or Data Corruption Reported

Issue: Fatal error(s) or data corruption(s) are reported when accessing virtual disks.

**Corrective
Action:** Contact Dell Technical Support.

Physical Disk Displayed As Blocked

Issue: One or more physical disks is displayed as **Blocked** and cannot be configured.

**Corrective
Action:** Update the PERC cards to the latest firmware available on dell.com/support.


Multiple Disks Become Inaccessible

Issue:	Rebuilding the physical disks after multiple disks become simultaneously inaccessible.
Corrective Action:	Multiple physical disk errors in a single array typically indicate a failure in cabling or connection and could involve the loss of data. You can recover the virtual disk after multiple physical disks become simultaneously inaccessible. Perform the following steps to recover the virtual disk:

 **CAUTION:** Follow the safety precautions to prevent electrostatic discharge.


1. Turn off the system, check cable connections, and reseal physical disks.
2. Ensure that all the disks are present in the enclosure.
3. Turn on the system and enter the **BIOS Configuration Utility** (<Ctrl> <R>) utility.
4. Import the foreign configuration.
5. Press <F> at the prompt to import the configuration, or press <C> to enter the **BIOS Configuration Utility** (<Ctrl> <R>) and either import or clear the foreign configuration.

If the virtual disk is redundant and transitioned to **Degraded** state before going **Offline**, a rebuild operation starts automatically after the configuration is imported. If the virtual disk has gone directly to the **Offline** state due to a cable pull or power loss situation, the virtual disk is imported in its **Optimal** state without a rebuild occurring.

 **NOTE:** You can use the **BIOS Configuration Utility** (<Ctrl> <R>) or Dell OpenManage storage management application to perform a manual rebuild of multiple physical disks. For information on rebuilding a single physical disk, see the topic [Performing A Manual Rebuild Of An Individual Physical Disk](#).

Rebuilding A Failed Physical Disk

Issue:	Rebuilding a physical disk after one of them is in a failed state.
Corrective Action:	If you have configured hot spares, the PERC card automatically tries to use one of them to rebuild a physical disk that is in a failed state. Manual rebuild is necessary if no hot spares with enough capacity to rebuild the failed physical disks are available. You must insert a physical disk with enough storage in the subsystem before rebuilding the physical disk.

 **NOTE:** You can use the **BIOS Configuration Utility** (<Ctrl> <R>) or Dell OpenManage storage management application to perform a manual rebuild of an individual physical disk. For information on rebuilding a single physical disk, see the topic [Performing A Manual Rebuild Of An Individual Physical Disk](#).

Virtual Disk Fails During Rebuild Using A Global Hot Spare

Issue:	A virtual disk fails during rebuild while using a global hot spare.
Description:	The global hot spare goes back to Hotspare state and the virtual disk goes to Failed state.

Virtual Disk Fails During Rebuild Using A Dedicated Hot Spare

Issue:	A virtual disk fails during rebuild while using a dedicated hot spare.
Description:	The dedicated hot spare goes to Ready state and the virtual disk goes to Failed state.

Physical Disk Fails During Reconstruction On Redundant Virtual Disk

Issue:	A physical disk fails during a reconstruction process on a redundant virtual disk that has a hot spare.
Corrective Action:	The dedicated hot spare goes to Ready state and the virtual disk goes to Failed state.

Virtual Disk Fails Rebuild Using A Dedicated Hot Spare


Issue:	A virtual disk fails during rebuild while using a dedicated hot spare.
Description:	The dedicated hot spare goes to Ready state and the virtual disk goes to Failed state.

Physical Disk Takes A Long Time To Rebuild



Issue:	A physical disk is taking longer than expected to rebuild.
Description:	A physical disk takes longer to rebuild when under high stress. For example, there is one rebuild I/O operation for every five host I/O operations. You cannot add a second virtual disk to a disk group while the virtual disk in that disk group is undergoing a rebuild. The firmware does not allow you to create a virtual disk using the free space available in a disk group if a physical disk in a virtual disk group is undergoing a rebuild operation.

SMART Errors


SMART monitors the internal performance of all motors, heads, and physical disk electronics and detects predictable physical disk failures.

 **NOTE:** For information about where to find reports of SMART errors that could indicate hardware failure, see the Dell OpenManage storage management documentation at dell.com/support/manuals.

Smart Error Detected On A Physical Disk In A Redundant Virtual Disk

Issue:	A SMART error is detected on a physical disk in a redundant virtual disk.
Corrective Action:	<p>Perform the following steps:</p> <ol style="list-style-type: none">1. Back up your data.2. Force the physical disk offline.  NOTE: If a hot spare is present, the rebuild starts with the hot spare after the disk is forced offline.3. Replace the disk with a new physical disk of equal or higher capacity.4. Perform the Replace Member operation.  NOTE: The Replace Member operation allows you to copy data from a source physical disk of a virtual disk to a target physical disk that is not a part of the virtual disk. For more information about the Replace Member feature, see the topic Using Replace Member And Revertible Hot Spares.

Smart Error Detected On A Physical Disk In A Non-Redundant Virtual Disk

Issue:	A SMART error is detected on a physical disk in a redundant virtual disk.
Corrective Action:	Perform the following steps: <ol style="list-style-type: none">1. Back up your data.2. Use Replace Member or set up a global hot spare to replace the disk automatically.<div> NOTE: For more information about the Replace Member feature, see the topic Using Replace Member And Reversible Hot Spares.</div>3. Replace the affected physical disk with a new physical disk of equal or higher capacity.4. Restore from the backup.

Replace Member Errors

 **NOTE:** For more information about the **Replace Member** features, see the topic [Using Replace Member And Reversible Hot Spares](#).

Source Disk Fails During Replace Member Operation

Issue:	The source disk fails during the Replace Member operation.
Corrective Action:	If the source data is available from other disks in the virtual disk, the rebuild begins automatically on the target disk, using the data from the other disks.

Target Disk Fails

Issue:	The target disk fails.
Corrective Action:	If the target disk fails, the Replace Member operation aborts.

General Disk Fails

Issue:	A general disk fails.
Corrective Action:	If the target disk fails and the Replace Member operation aborts but the source data is still available, then the Replace Member operation continues as Replace Member .

Linux Operating System Errors

Virtual Disk Policy Is Assumed As Write-Through Error Message

Error:	<pre><Date:Time> <HostName> kernel: sdb: asking for cache data failed <Date:Time> <HostName> kernel: sdb: assuming drive cache: write through</pre>
---------------	---

Corrective Action: The error message is displayed when the Linux Small Computer System Interface (SCSI) mid-layer asks for physical disk cache settings. The controller firmware manages the virtual disk cache settings on a per controller and a per virtual disk basis, so the firmware does not respond to this command. The Linux SCSI mid-layer assumes that the virtual disk's cache policy is **Write-Through**. SDB is the device node for a virtual disk. This value changes for each virtual disk.

See the topic [Write-Back And Write-Through](#) for more information about **Write-Through** cache. Except for this message, there is no effect of this behavior on normal operation. The cache policy of the virtual disk and the I/O throughput are not affected by this message. The cache policy settings for the PERC SAS RAID system remain unchanged.

Driver Does Not Auto-Build Into New Kernel

Issue: Driver does not auto-build into new kernel after customer updates.

Corrective Action: This error is a generic problem for Dynamic Kernel Module Support (DKMS) and applies to all DKMS-enabled driver packages. This issue occurs when you perform the following steps:

1. Install a DKMS-enabled driver package.
2. Run **up2date** or a similar tool to upgrade the kernel to the latest version.
3. Reboot to the new kernel.

The driver running in the new kernel is the native driver of the new kernel. The driver package you installed previously in the new kernel does not take effect in the new kernel:

1. **Type:** `dkms build -m <module_name> -v <module version> -k <kernel version>.`
2. **Type:** `dkms install -m <module_name> -v <module version> -k <kernel version>.`
3. **Type** the following to check whether the driver is successfully installed in the new kernel: DKMS. The following details appear:
`<driver name>, <driver version>, <new kernel version>:
installed`

Unable To Register SCSI Device Error Message

Error: `smartd[smartd[2338] Device: /dev/sda, Bad IEC (SMART) mode page, err=-5, skip device smartd[2338] Unable to register SCSI device /dev/sda at line 1 of file /etc/smartd.conf.`

Corrective Action: This is a known issue. An unsupported command is entered through the user application. User applications attempt to direct Command Descriptor Blocks to RAID volumes. The error message does not affect the feature functionality. The **Mode Sense/Select** command is supported by firmware on the controller. However, the Linux kernel **daemon** issues the command to the virtual disk instead of to the driver **IOCTL** node. This action is not supported.

Disk Carrier LED Indicators

The LED on the physical disk carrier indicates the state of each physical disk. Each disk carrier in your enclosure has two LEDs: an activity LED (green) and a status LED (bicolor, green/amber). The activity LED is active whenever a disk is being accessed while the status LED indicates when a disk is being spun up, is rebuilding, or is in a fault state.




Figure 5. Disk Carrier LED Indicators

1. activity LED
2. status LED

Appendix: RAID Description


RAID is a group of independent physical disks that provides high performance by increasing the number of disks used for saving and accessing data.

 **CAUTION:** In the event of a physical disk failure, a RAID 0 virtual disk fails, resulting in data loss.

A RAID disk subsystem offers the following benefits:

- Improved I/O performance and data availability.
- Improved data throughput because several disks are accessed simultaneously. The physical disk group appears either as a single storage unit or multiple logical units to the host system.
- Improved data storage availability and fault tolerance. Data loss caused by a physical disk failure can be recovered by rebuilding missing data from the remaining physical disks containing data or parity.

Summary Of RAID Levels

 **NOTE:** PERC H710, H710P and H810 cards support all RAID levels listed below. PERC H310 supports RAID 5 with limited performance and does not support RAID 6 and RAID 60.

- RAID 0 uses disk striping to provide high data throughput, especially for large files in an environment that requires no data redundancy.
- RAID 1 uses disk mirroring so that data written to one physical disk is simultaneously written to another physical disk. RAID 1 is good for small databases or other applications that require small capacity and complete data redundancy.
- RAID 5 uses disk striping and parity data across all physical disks (distributed parity) to provide high data throughput and data redundancy, especially for small random access.
- RAID 6 is an extension of RAID 5 and uses an additional parity block. RAID 6 uses block-level striping with two parity blocks distributed across all member disks. RAID 6 provides protection against double disk failures, and failures while a single disk is rebuilding. If you are using only one array, deploying RAID 6 is more effective than deploying a hot spare disk.
- RAID 10 is a combination of RAID 0 and RAID 1, uses disk striping across mirrored disks. It provides high data throughput and complete data redundancy. RAID 10 can support up to eight spans, and up to 32 physical disks per span.
- RAID 50 is a combination of RAID 0 and RAID 5 where a RAID 0 array is striped across RAID 5 elements. RAID 50 requires at least six disks.
- RAID 60 is a combination of RAID 0 and RAID 6 where a RAID 0 array is striped across RAID 6 elements. RAID 60 requires at least eight disks.

RAID Terminology

Disk Striping

Disk striping allows you to write data across multiple physical disks instead of just one physical disk. Disk striping involves partitioning each physical disk storage space in stripes of the following sizes: 64 KB, 128 KB, 256 KB, 512 KB, and 1024 KB. The stripes are interleaved in a repeated sequential manner. The part of the stripe on a single physical disk is called a stripe element.

For example, in a four-disk system using only disk striping (used in RAID 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple physical disks are accessed simultaneously, but disk striping does not provide data redundancy.





Figure 6. Example of Disk Striping (RAID 0)

Disk Mirroring

With mirroring (used in RAID 1), data written to one disk is simultaneously written to another disk. If one disk fails, the contents of the other disk can be used to run the system and rebuild the failed physical disk. The primary advantage of disk mirroring is that it provides complete data redundancy. Both disks contain the same data at all times. Either of the physical disks can act as the operational physical disk.

Disk mirroring provides complete redundancy, but is an expensive option because each physical disk in the system must be duplicated.

 **NOTE:** Mirrored physical disks improve read performance by read load balance.




Figure 7. Example of Disk Mirroring (RAID 1)

Spanned RAID Levels

Spanning is a term used to describe the way in which RAID levels 10, 50, and 60 are constructed from multiple sets of basic, or simple RAID levels. For example, a RAID 10 has multiple sets of RAID 1 arrays where each RAID 1 set is considered a span. Data is then striped (RAID 0) across the RAID 1 spans to create a RAID 10 virtual disk. Similarly, RAID 50 and RAID 60 combine multiple sets of RAID 5 or RAID 6 respectively with striping.

Parity Data

Parity data is redundant data that is generated to provide fault tolerance within certain RAID levels. In the event of a disk failure, the parity data can be used by the controller to regenerate user data. Parity data is present for RAID 5, 6, 50, and 60.

The parity data is distributed across all the physical disks in the system. If a single physical disk fails, it can be rebuilt from the parity and the data on the remaining physical disks. RAID level 5 combines distributed parity with disk striping. Parity provides redundancy for one physical disk failure without duplicating the contents of the entire physical disks.

RAID 6 combines dual distributed parity with disk striping. This level of parity allows for two disk failures without duplicating the contents of entire physical disks.





Figure 8. Example of Distributed Parity (RAID 5)

 **NOTE:** Parity is distributed across multiple physical disks in the disk group.





Figure 9. Example of Dual Distributed Parity (RAID 6)

 **NOTE:** Parity is distributed across all disks in the array.